ESMDynamic: Fast and Accurate Prediction of Protein Dynamic Contact Maps from Single Sequences

Diego E. Kleiman,[†] Jiangyan Feng,[‡] Zhengyuan Xue,[†] and Diwakar Shukla[*,†,‡,¶,§,‖]

†*Center for Biophysics and Quantitative Biology, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*

‡*Department of Chemical and Biomolecular Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*

¶*Department of Chemistry, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*

§*Department of Bioengineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*

‖*Department of Plant Biology, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA*

E-mail: diwakar@illinois.edu

Abstract

Understanding conformational dynamics is essential for elucidating protein function, yet most deep learning models in structural biology predict only static structures. Here, we introduce ESMDynamic, a deep learning model that predicts dynamic residue-residue contact probability maps directly from protein sequence. Built on the ESMFold architecture, ESMDynamic is trained on contact fluctuations from experimental structure ensembles and molecular dynamics (MD) simulations, enabling it to capture diverse modes of structural variability without requiring multiple sequence alignments. We benchmark ESMDynamic on two large-scale MD datasets (mdCATH and ATLAS), showing that it matches or outperforms state-of-the-art ensemble prediction models (AlphaFlow, ESMFlow, BioEmu) for transient contact prediction while offering orders-of-magnitude faster inference. We demonstrate the model on the ASCT2 and SWEET2b transporters, a *de novo* troponin C design, and the HIV-1 protease homodimer, illustrating generalization to unseen systems and recovery of experimentally validated dynamic contacts. Furthermore, we present an automated pipeline using ESMDynamic predictions to select collective variables for Markov State Model construction, producing high-quality kinetic models from unbiased MD simulations of SWEET2b. Overall, ESMDynamic provides a compact and interpretable sequence-based description of conformational dynamics, with broad applications in protein engineering, functional analysis, and simulation-guided discovery.

Introduction

In recent years, machine-learning (ML) models have achieved excellent accuracy in predicting protein structure from amino acid sequences by exploiting their encoded evolutionary information. While some models[1–5] utilize multiple-sequence alignments (MSAs)[6,7] to acquire this information, other models employ protein language models (PLMs)[8–10] to embed sequences in a latent space whose features facilitate downstream tasks, including structure prediction. Proteins always exist as conformational ensembles in thermodynamic equilibrium.[11] While

knowledge of protein structure is crucial for understanding function, the current paradigm establishes the role of dynamics as a necessary hinge point between structure and function,[12] a relationship that was determined as early as the first cooperative binding models[13] and was later verified in many instances of allosteric regulation,[14] motor proteins,[15] and enzymatic catalysis,[16] among others.[17–20] However, structure prediction models only output static structures. Even in cases where a sequence maps to different stable conformations,[21] these models tend to suffer from mode collapse,[22] where a single high-likelihood prediction is favored over other possible outputs. Nonetheless, several methodological innovations made it possible to predict multiple structures from sequences.[23–27]

Due to the relevance of dynamics in the elucidation of protein function, many methods have been devised to predict structural ensembles. Many such protocols rely on partially hiding, masking, or noising the coevolutionary data passed to the models in a random[23] or a biased[24] fashion to introduce diversity among the predictions. However, these methods present certain shortcomings. Firstly, the model must be iteratively queried with the perturbed inputs, increasing the inference time. Secondly, the sensitivity of the models to input perturbations is poorly characterized. Therefore, it cannot be known *a priori* whether alternative conformations will be predicted, and if they are, it is challenging to design the perturbation that results in the desired output change. Some authors have suggested that this approach is successful only when the alternative conformations are observed during training,[22] although results with held-out test sets indicate the contrary.[3] These methods have been extended by incorporating molecular dynamics (MD) simulations to sample a free energy landscape of the system once a conformational ensemble has been predicted.[28] Another family of methods employs diffusion-based or flow-based probabilistic models to sample conformational ensembles.[26,27]

A different approach to predict alternative conformations relies on re-interpreting the output of structure prediction models to infer the implicit conformational diversity encoded within intermediate outputs, which is typically suppressed once the model converges on a fi-

nal structure. These approaches may construct energy restrictions based on pairwise residue distances.[29] However, these approaches are only intended to work on capturing large conformational changes, mostly rearrangements of interdomain distances. Nonetheless, many functionally relevant transition are better described as subtle fluctuations propagated through networks of dynamically-coupled residues.[30]

With this in mind, we propose to leverage structure prediction models to predict dynamic contact maps of proteins. Structure prediction models introduce a strong inductive bias to learn protein dynamics. This arises from the fact that structural knowledge is required to propose physically coherent hypotheses about the dynamics of the system. Moreover, the evolutionary signals that encode structural information are also likely to encode the dynamics.[10]

While dynamic contact maps represent a coarse view of protein dynamics, many collective variables (CVs) used to represent conformational landscapes in proteins are purely functions of inter-residue distances, since these pairwise distances define the geometry of the backbone. Understanding the dynamic contacts within proteins can yield valuable insights for both experimental and computational studies. For instance, when designing Förster Resonance Energy Transfer (FRET) or Double Electron-Electron Resonance (DEER) experiments, one must carefully select labeling sites that undergo dynamic changes to effectively capture the conformational ensemble of the system.[31,32] Similarly, mutational studies can benefit from targeting residues involved in dynamic contacts, as these often represent key functional or regulatory regions.[33,34] In computational research, such as MD simulations, the sheer number of particles makes it challenging to monitor all degrees of freedom. Dynamic contact maps can help define meaningful CVs for enhanced sampling protocols, where an external bias is applied along the CVs,[35] or for efficient *post hoc* featurization and dimensionality reduction of trajectory data.[36]

In this work, we posit that dynamics can be learned by finetuning structure prediction models, and introduce ESMDynamic, a sequence to dynamic contact map model. We employ

ESMFold[8] as a base model, which is based on ESM2, a PLM that extracts coevolutionary signals without requiring a MSA, and can therefore infer structures without employing expensive sequence alignments, unlike other established methods.[1,3–5] A two-stage training protocol was applied to fit ESMDynamic, where we first pretrained on experimentally resolved structures and AlphaFold DB[37] models from 88,294 sequence-aligned clusters ($> 95\%$ intracluster identity), where structural variations were used to infer dynamic contacts. These clusters comprise 946,397 protein chains from 598,144 uniquely resolved structures, with a median of four conformers per protein. The model was then finetuned on mdCATH,[38] a dataset of MD simulations covering 5,398 soluble domains, with an average trajectory length of 464 ns. The training objective was to classify protein residue-residue pairs as dynamic contacts, labeling non-interacting or static pairs as negatives and fluctuating pairs as positives.

Overall, ESMDynamic achieves excellent performance on the sequence-identity aware test split ($< 20\%$ identity between clusters in train, validation, and test sets), with a balanced accuracy of 80% and recall of 73% when using MD simulations as ground truth. Notably, more than 90% of all positive predictions occur within five residues of true positives, highlighting a strong correlation with the underlying dynamics. Importantly, ESMDynamic surpasses state-of-the-art ensemble prediction models[27,39] in the identification of dynamic contacts, both in performance metrics and inference wall-clock time. This makes ESMDynamic a valuable tool for rapidly predicting dynamic properties in seconds, inlcuding features that may be utilized as CVs, obtaining insights that would take weeks or months with MD simulations.

Results

Model architecture

In this work, we define dynamic contacts as residue pairs that transiently come into spatial proximity using a customary threshold of 8 Å.[40] While dynamic contacts frequently overlap with native contacts, the former are not necessarily a subset of the latter; residue pairs which form short-lived or unstable interactions are also considered dynamic contacts. Moreover, the definition of dynamic contacts involves nuances which are often ignored when defining native contacts. While native contacts are generally defined based on a single static structure, dynamic contacts reflect the intrinsic flexibility of proteins and require considering an ensemble of conformations. In principle, a rigorous definition of dynamic contacts would depend on thermodynamic parameters, such as temperature, that determine the permissible states within the ensemble.[41] However, in practice, defining such ensembles is challenging due to variability in the experimental conditions under which protein structures are determined. To address this, we adopt a data-driven approach: for proteins with multiple experimentally determined structures, we identify dynamic contacts as residue pairs whose spatial separation varies above and below the threshold across conformations. A similar approach is applied to MD simulation datasets, where contact variability over time serves as the basis for labeling (see Methods). This framework allows us to systematically annotate dynamic contacts across diverse datasets. ESMDynamic is trained to predict these dynamic contact maps, where each element of the matrix represents the probability that a given residue pair meets the dynamic contact criteria.

ESMDynamic (Figure 1) is built upon the ESMFold framework. Structure prediction models offer a powerful prior for inferring protein dynamics, as knowledge of 3D geometry constrains the plausible ways a protein conformation can fluctuate. By leveraging ESMFold, we incorporate structural reasoning directly into the prediction process. An additional advantage of ESMFold is that it extracts evolutionary and structural features from a sin-

gle input sequence, removing the need for computationally expensive MSAs and enabling high-throughput predictions.

Beyond its core structure prediction capability, ESMFold provides auxiliary signals that are particularly informative for modeling conformational variability. These include distogram logits, which represent predicted distributions of residue-residue distances, and can highlight residue pairs that fluctuate between contact and non-contact states.[1] Language model (LM) logits provide token-level probabilities that correlate with sequence conservation, information often enriched at functionally or structurally significant sites.[8] Uncertainty-related outputs, such as predicted TM-score (pTM) and per-residue confidence (pLDDT), further help distinguish between rigid and flexible regions of a protein.[1] ESMDynamic integrates these auxiliary signals by biasing the final sequence and pairwise representations from ESMFold's last Evoformer[1] block. These enriched features are then passed to a dedicated *Dynamic Contact Module* (DCM), composed of two additional Evoformer blocks. The Evoformer architecture is well-suited for this task, as it models complex, context-dependent relationships between residues using attention-based mechanisms. Table S1 provides a detailed breakdown of the layers within each submodule, along with their respective parameter counts.

As mentioned, the model outputs a probabilistic contact map, reflecting the likelihood of each residue pair forming a dynamic interaction. To focus learning on difficult examples, we use the focal loss function,[42] which emphasizes instances with high probability of a false prediction by dynamically reweighting the loss contribution from each sample during training.

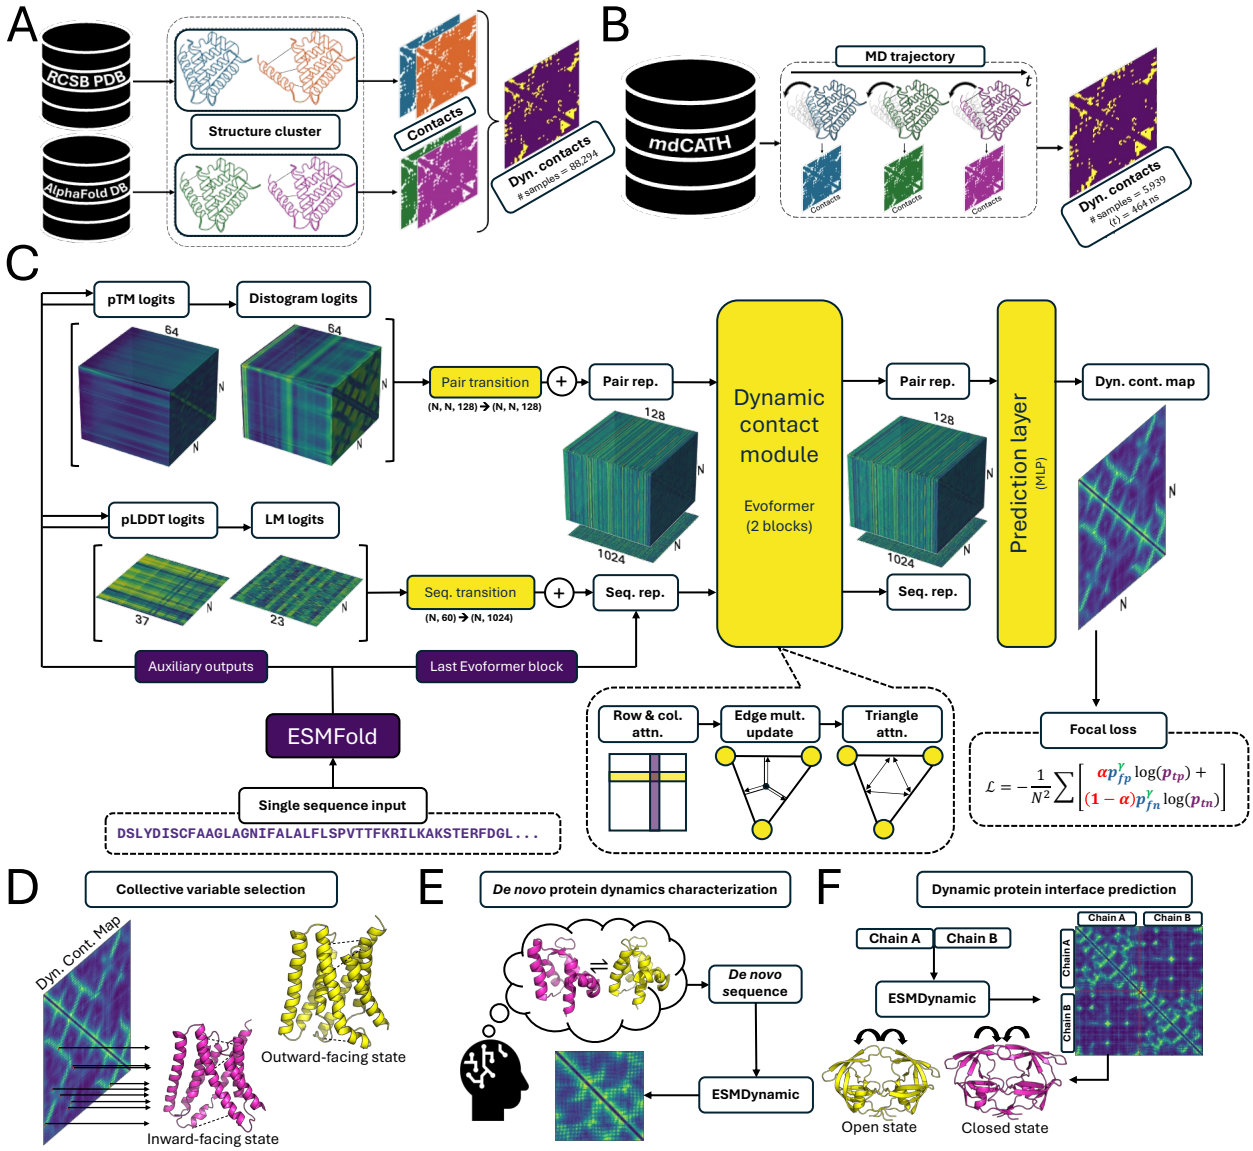

Figure 1: Schematic of the ESMDynamic approach. (A) Sequence-aligned structure data sets were employed to generate dynamic contact maps for pretraining. (B) Dynamic contacts extracted from MD trajectories in mdCATH are used in training. (C) Architecture schematic: a single input sequence is passed through ESMFold. Auxiliary outputs are processed by transition modules composed of linear and normalization layers. These features are used to bias the output from the last Evoformer block of ESMFold, which is passed to the Dynamic Contact Module (DCM). Finally, a multilayer perceptron (MLP) converts the pair representation into dynamic contact probabilities. The focal loss is employed ($p_{(fp/fn)}$: probability of false positive/negative, $p_{(tp/tn)}$: probability of true positive/negative, α: positive class weight, γ: difficulty boost). The resulting dynamic contact map can be used for numerous downstream analyses. (D) Selected residue-residue distances can be used as collective variables for MD trajectory analysis or biasing. (E) ESMDynamic can be used to predict dynamic contacts in *de novo* proteins. (F) Emergent properties of ESMDynamic allow it to detect dynamic protein interfaces.

Performance and ablation studies

ESMDynamic was first pretrained on experimentally determined protein structures and subsequently fine-tuned using MD simulations (see Methods). Our validation focuses primarily on the MD dataset. MD trajectories offer a continuous and detailed view of protein conformational variability, whereas experimental structures typically provide only sparse snapshots, limiting their utility for capturing dynamic behavior. This is reflected in the proportion of dynamic contacts resolved by each method; in experimental clusters, dynamic contacts constitute fewer than 1% of all residue pairs, while in MD simulations, this proportion increases to approximately 15%, reflecting their greater sensitivity to conformational dynamics. On the MD-based, sequence-identity aware test split, ESMDynamic achieves a balanced accuracy of 80% (Figure 2A) and a recall of 73%.

To assess the contribution of different model components, we also evaluated ablated variants lacking auxiliary inputs, the DCM, or the final prediction layer (see Methods). As expected, the removal of the model elements affects its final performance across the training, validation, and test sets (Figure 2A). While the first two reduced models show marginal degradation in performance, the ablation lacking a prediction head shows a significant degradation in balanced accuracy. The ablation results highlight the fact that the features derived from ESMFold can be readily used to predict dynamic contacts, i.e., the model without specialized Evoformer blocks can still attain good performance. However, the prediction task is complex and nuanced: the addition of auxiliary features and a dedicated DCM does improve the performance, indicating the need for additional parameters to correctly predict difficult examples. On the other hand, a model that utilizes a simple pooling operation on the pair representation instead of a prediction head suffers a large decline in performance.

Beyond conventional classification metrics, we also evaluate the spatial correlation between predicted and ground-truth dynamic contacts in sequence space. Specifically, we compute the probability distribution of the shortest L_1 distance between predicted and true positives (Figure 2C). This distribution's y-axis intercept is related to precision, and its over-

all shape captures how tightly the model's predictions cluster around true dynamic contacts. Figure S1 shows an illustration of how the L_1 distance is calculated. We find that $> 90\%$ of positive predictions fall within five residues of a true positive. To quantify this further, we fit an exponential function to the distribution and define its characteristic decay length as a proxy for the effective resolution of the model. This value represents the radius within which the predictive signal tends to concentrate around ground-truth positives, with a result of 1.29 residues for the full model.

Since we are computing the shortest L_1 distance to characterize the resolution of the model, we note that even a naive classifier would show strong results if ground-truth samples were distributed roughly uniformly throughout the map. However, this is not the case, as shown by the probability distribution for the naive classifier (Figure 2C). To further assess the co-localization of dynamic contacts, we compute their L_1 radial distribution function (Figure S2), which confirms that dynamic contacts are spatially correlated in sequence space and that the concentration of positive predictions around the ground-truth is a feature of such correlation. Ablated variants of the model display slightly poorer resolution, as shown by the slower decay in the probability distribution (i.e., longer characteristic lengths). The y-axis cutoff is also lower for the ablated variants, indicating that precision is marginally degraded when removing different model parts. Figure S3 shows the full exponential fit on ablated models.

ESMDynamic surpasses ensemble prediction models in accuracy and runtime

We compared ESMDynamic to three state-of-the-art ensemble prediction models: BioEmu,[39] AlphaFlow, and ESMFlow.[27] While BioEmu and AlphaFlow are both based on AlphaFold2 and rely on MSAs, ESMFlow is built on the ESMFold architecture. Unlike ESMDynamic, which is specifically trained to predict dynamic contacts, these models are designed to generate full structural ensembles from sequence inputs. They learn to reconstruct conforma-

tional variability using both experimental structures and MD simulations. In particular, BioEmu was trained on approximately 200 ms of MD data,[39] whereas the flow models and ESMDynamic were trained on only a few milliseconds of protein trajectories. Although ensemble-generation models are not explicitly trained to classify dynamic contacts, such contacts can be inferred from their predictions, as dynamic contacts are inherent features of structural variability. The models compared here serve as demanding benchmarks, as they have demonstrated superior performance over other ensemble-generation methods, including those based on MSA subsampling.[27,39]

When evaluated on dynamic contact prediction, ESMDynamic matches or substantially outperforms all benchmarks (Figure 2B). ESMDynamic and BioEmu achieve similar balanced accuracy on the mdCATH test set. A large portion (46 ms) of BioEmu's training data was derived from CATH sequences,[39] so overlap between the test set used in this work and BioEmu's training set cannot be ruled out. In contrast, the flow models tend toward high-precision, low-recall predictions, resulting in lower overall balanced accuracy. Specifically, on the mdCATH test set, AlphaFlow achieved the highest precision (84%), but its recall was low (40%), yielding a balanced accuracy of 67% averaged across proteins. All models were also evaluated on the AlphaFlow test split derived from ATLAS.[43] Although mdCATH and ATLAS are both MD datasets, they differ in key parameters, including force field choice and simulation temperature.[38,43] Notably, the flow models were trained on the ATLAS ensembles, whereas ESMDynamic had never seen this data during training. Despite this, ESMDynamic significantly outperforms the flow models and BioEmu in identifying dynamic contacts, achieving a balanced accuracy of 88% (Figure 2B). While BioEmu was not trained on ATLAS data, it performs comparably to the best flow model, AlphaFlow.

The underperformance of the flow models is likely due to their design, which samples from a distribution constrained by the AlphaFold2 or ESMFold structure module (i.e., noised inputs are always passed through the structure module to generate the final conformation, rather than using freshly trained weights). These constraints favor the generation of plau-

sible, low-energy conformations but limit the models' ability to represent low-probability states in which contacts are broken. In contrast, ESMDynamic is trained directly to predict dynamic contact probabilities without imposing strict geometric consistency. As a result, it can assign high probabilities to multiple, possibly incompatible contact states, better reflecting the overall potential fluctuations in protein structures. We also compared a naive baseline that labels the native contacts predicted by ESMFold as dynamic. While this baseline performs similarly to the distilled version of ESMFlow on mdCATH, the true dynamic landscape is more complex than what is revealed by native contacts, which explains its poor accuracy.

In terms of runtime, ESMDynamic predicts dynamic contact maps in just a few seconds to a few minutes, depending on protein length, since it requires only a single forward pass. Its runtime is virtually identical to that of ESMFold. In contrast, generative methods require up to hundreds of forward passes per protein to sample enough structures, significantly increasing computation time. Among the generative models compared, the distilled version of ESMFlow was the fastest, albeit at the cost of lower accuracy. For example, in our evaluation on the ATLAS test set (82 proteins), ESMDynamic completed predictions in 9.45 minutes on a single NVIDIA RTX 4090 GPU, while the fastest generative method took approximately 7.9 hours using 250 samples per protein, as in the original ESMFlow protocol.[27]

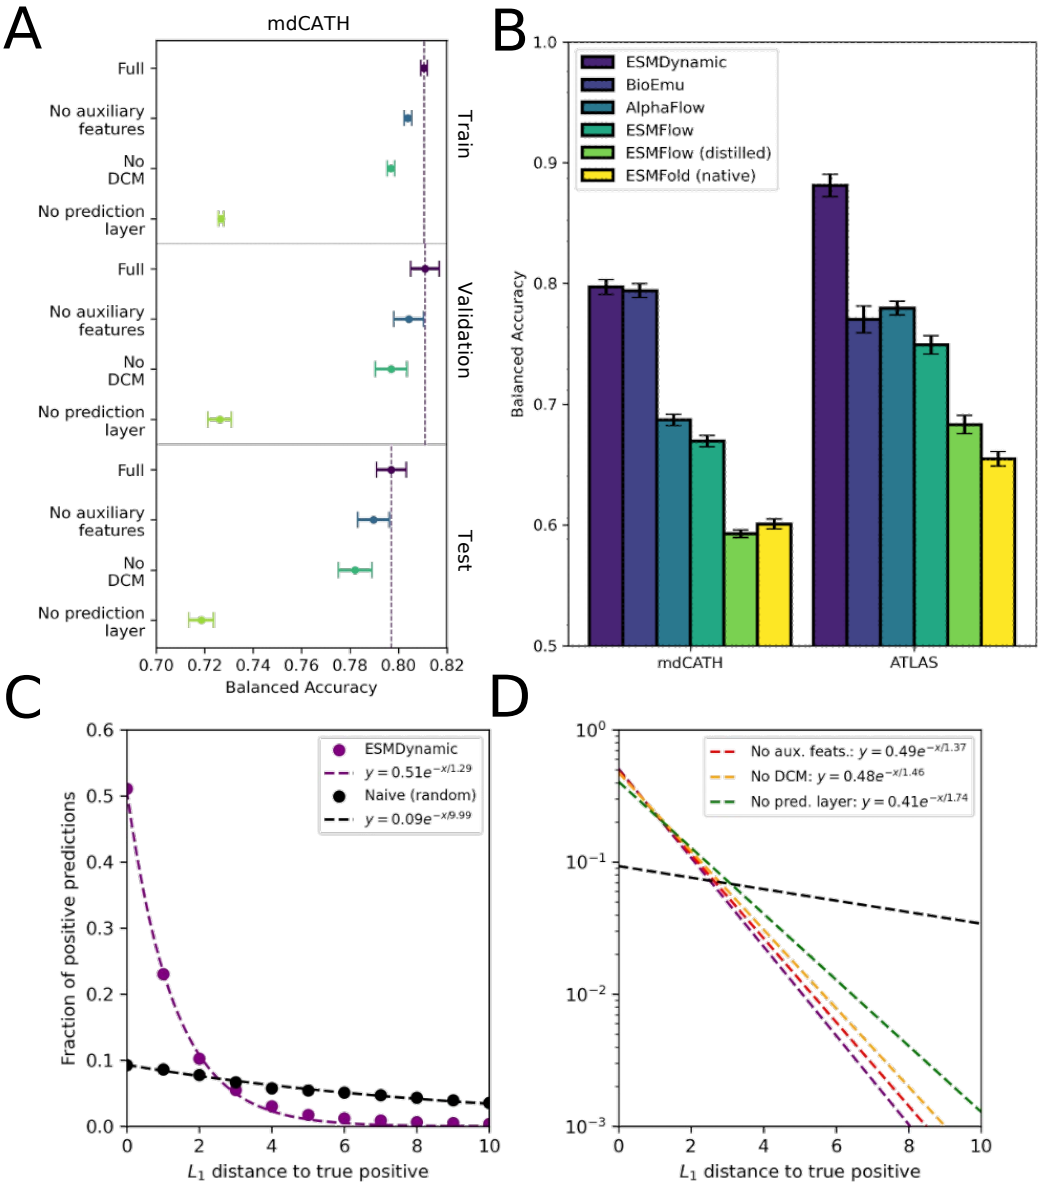

Figure 2: Performance of ESMDynamic, ablated variants, and comparison with other methods. (A) Balanced accuracy of the full model and ablated variants across mdCATH test splits. Error bars are standard errors. (B) Comparative performance of ESMDynamic, BioEmu, AlphaFlow, and ESMFlow models on the mdCATH and ATLAS test sets. Error bars are standard errors. Naive baseline (ESMFold) assumes all native contacts are dynamic. (C) Fraction of positive predictions as a function of sequence-space distance to ground-truth positives. ESMDynamic exhibits an exponential decay, indicating that predicted positives are highly localized around ground-truth positives, in contrast to a random naive model that produces a more uniform distribution. (D) The exponential fit from (C) on a logarithmic scale highlights differences in decay rates, with the full model showing the steepest negative slope relative to ablated variants and the naive baseline.

In addition to evaluating our model against ablated variants and generative models, we also examine the correlation between ESMDynamic's predictions and the unsupervised contact probabilities produced by ESM2.[8] Similar to the comparison against the naive baseline in Figure 2B, this test serves to assess whether ESMDynamic is simply recovering native contact patterns rather than learning features specific to conformational dynamics. If the model's output were highly correlated with ESM2's unsupervised contact predictions, it would suggest that the dynamic contact probabilities reflect static structural features rather than dynamic behavior. However, on the mdCATH test set, we observe no meaningful correlation between ESMDynamic's predictions and those from ESM2 (Figure S4). This supports the conclusion that ESMDynamic is learning non-trivial, dynamics-specific signals that are distinct from those captured by unsupervised models trained solely on sequence data.

ESMDynamic recovers experimental dynamics in proteins beyond its training set

Figure 3 illustrates the application of ESMDynamic to the neutral amino acid transporter ASCT2/SLC1A5, which is upregulated in cancer and also serves as a retroviral receptor.[44] The ASCT2 sequence shares less than 20% identity with proteins in both the pretraining and finetuning datasets (see Methods). Although ESMDynamic was pretrained on diverse structural ensembles, it was subsequently finetuned on the mdCATH database, which lacks membrane proteins, raising the question of whether the model can retain accuracy on transmembrane transporters. These factors make ASCT2 a challenging test case.

To evaluate performance, we analyzed two experimental structures corresponding to inward-facing (PDB 6GCT)[45] and outward-facing (PDB 7BCQ)[46] conformations, which define a set of experimentally validated dynamic contacts (highlighted with red circles in Figure 3, Right). ESMDynamic accurately recapitulates all known dynamic contacts and additionally predicts novel contacts in other regions of the protein. For instance, the model identifies a discrete dynamic contact cluster (Figure 3, Right inset) between residues S351–

A355 of hairpin 1 (HP1) and G430–P432 of hairpin 2 (HP2). This interface is important as HP2 constitutes both the extracellular and intracellular gate[47] and must move to admit or expel the substrate. Based on experimental structures, the interface between HP1 and HP2 is disrupted in the *apo* outward-facing state (PDB 6MP6[48]) and when the proline-derived inhibitor Lc-BPE binds at S353 (PDB 7BCQ),[46] locking the transporter in the outward-facing conformation.[46] On the other hand, this interface is preserved in the inward-facing conformation (PDB 6GCT).[45]

In summary, ESMDynamic successfully detected key dynamic features governing the transition between the inward- and outward-facing states of a transmembrane transporter with minimal sequence similarity to any protein in its training data. The ability of the model to generalize to membrane proteins, despite their absence in mdCATH, suggests that the finetuning procedure did not lead to performance degradation in this class of proteins. It is important to note that, while ESMDynamic itself was not trained or pretrained on this transporter, ESMFold was likely trained on homologous sequences. Thus, the model benefits from a solid structural foundation in predicting dynamic contacts, even though the dynamics were never directly seen during training. Finally, the case of ASCT2 illustrates the potential of ESMDynamic to probe the dynamics of clinically relevant proteins, offering a scalable way to generate hypotheses about mechanisms of gating, inhibition, and transport. More broadly, these findings suggest that ESMDynamic can provide meaningful insights into protein classes beyond its training distribution, thereby extending its applicability to diverse problems in structural biology and drug discovery.

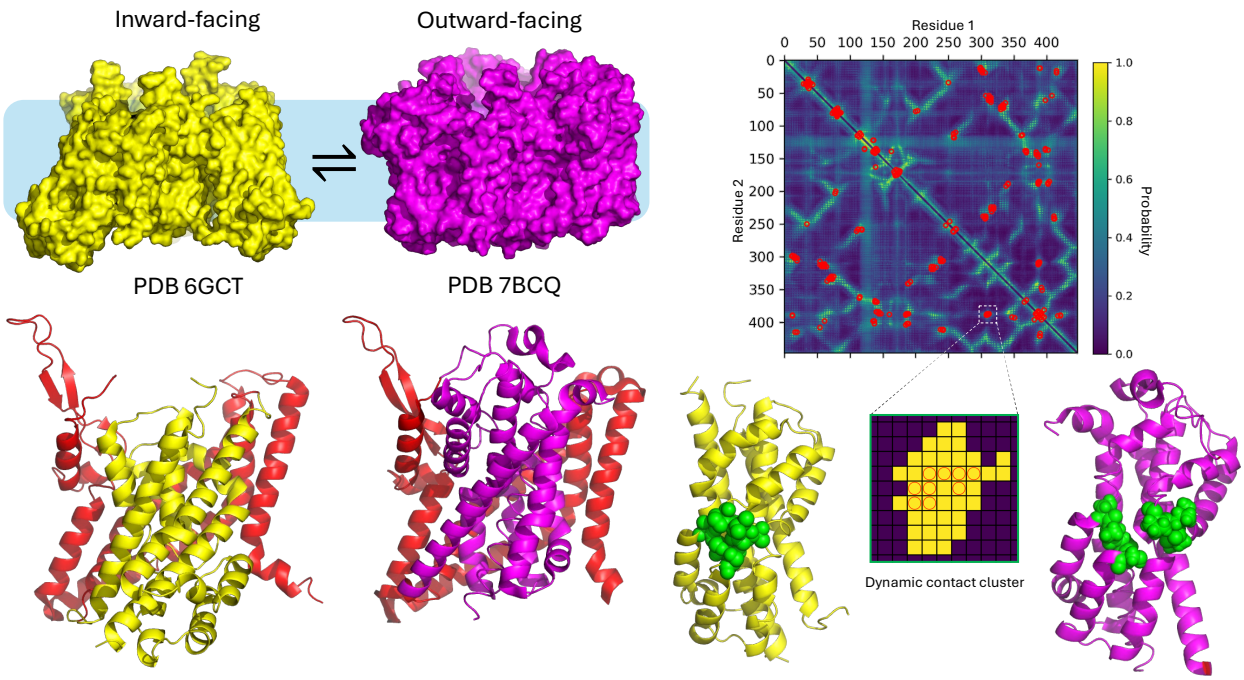

Figure 3: Sample output from ESMDynamic for the ASCT2/SLC1A5 transporter. Left: Experimental structures of the inward-facing and outward-facing conformational states. A cartoon representation highlights a structurally invariant region (red) and regions that undergo conformational change (yellow and purple, corresponding to each state). Right: Dynamic contact map predicted by ESMDynamic (residue indexing begins at D43). Red circles denote dynamic contacts inferred from the experimental structures, all of which are recovered by the model. The inset shows a zoomed-in view of a dynamic contact cluster involving residues S351 to A355 and G430 to P432. These contacts are present in the inward-facing state but are disrupted during the transition to the outward-facing state. Residues involved in these interactions are highlighted in green.

Dynamic contact maps enable automatic collective variable selection for MD simulations

A key application of ESMDynamic is leveraging the predicted dynamic contact map to identify relevant CVs for analysis or biasing in MD simulations. In this task, users examine the predicted dynamic contacts and employ either automated tools or knowledge-based criteria to select informative metrics to monitor. While the contact map highlights residue pairs expected to undergo distance changes, its utility extends beyond distance-based features. For instance, an interaction might break due to torsional changes within residues, and clusters

of dynamic contacts may signal larger conformational rearrangements.

In this section, we demonstrate an automated pipeline for CV selection to analyze 145 μs of unbiased MD simulations of *Oryza sativa* (rice) SWEET2b,[49] a transporter that facilitates passive sugar transport along concentration gradients with relevance to crop yield.[50] Notably, this protein was excluded from the ESMDynamic training set, with no homologs above 20% sequence identity present during training. This ensures the analysis reflects a realistic use case (i.e., applying ESMDynamic to a previously unseen protein sequence lacking close homologs in the training data).

The presented pipeline does not prescribe a fixed method for utilizing ESMDynamic predictions. Rather, it illustrates that a simple, fast analysis of predicted dynamic contacts can be highly effective in identifying meaningful CVs for MD trajectory analysis. As shown in Figure 4A, the pipeline begins by filtering residue-residue pairs to retain only high-order dynamic contacts. Specifically, predicted positive contacts between residues separated by at least 40 positions in sequence space are chosen. This step removes local dynamic contacts (e.g., those within transmembrane helices or between adjacent helices).

Subsequently, hierarchical clustering is applied to the remaining residue pairs to identify distinct dynamic regions within the protein. We treat the contact map positions as spatial "coordinates" and use single-linkage clustering with L_1 distances via Scipy,[51] limiting the maximum number of clusters to 9. This cap allows for a fair comparison between CVs selected via ESMDynamic and those chosen manually by a human expert, as described in ref. 49.

From each cluster, we select the highest-probability dynamic contact and track the corresponding inter-residue distances as CVs. These distances are then used to featurize the MD trajectory, followed by time-lagged independent component analysis (tICA)[52,53] for dimensionality reduction. tICA identifies the slowest decorrelating collective motions in the trajectory by maximizing autocorrelation over a specified lag time. This provides a reduced set of kinetically relevant coordinates that are particularly well suited for constructing Markov

State Models (MSMs). An MSM represents the conformational space as a set of discrete states and estimates the probabilities of transitioning between them over a given lag time. This framework enables the reconstruction of long-timescale dynamics from shorter MD trajectories by exploiting the Markovian property of memoryless transitions.[54] Following this framework, an MSM is fitted from the tICs using the same parameters as in ref. 49 (see also Supplementary Methods).

Figure 4B shows that the free energy landscape projected onto the top two tICA components successfully separates the transporter's functional states (the inward-facing and outward-facing conformations) displayed in yellow and purple, respectively, in Figures 4B,C. Furthermore, the MSM fitted on ESMDynamic-selected CVs achieves a VAMP-2 score[55] of 87% of the theoretical maximum, comparable to that obtained with human expert selection of CVs, which were based on visual inspection of trajectories to identify key interactions during state transitions.[49] In other words, ESMDynamic was able to recover CVs that are normally obtainable only through *post hoc* analysis of MD trajectories by relying solely on a protein sequence and a straightforward processing pipeline.

Figure 4C (center) shows the selected CVs as red lines projected onto both conformational states. ESMDynamic assigns high dynamic contact probabilities to residue pairs that act as intracellular and extracellular gates (F42–F164 and Y60–F191, respectively), which are highlighted as space-filling representations in Figure 4C (top and bottom). These gating interactions constitute central dynamic features of SWEET2b, as sugar translocation proceeds through an alternating-access mechanism in which one gate must open to permit substrate binding while the opposite gate subsequently opens to enable substrate release on the other side of the membrane. This example demonstrates that a trivial pipeline, starting from sequence-based dynamic contact predictions from ESMDynamic, can provide a rational framework for selecting CVs to investigate the mechanisms of complex transport proteins.

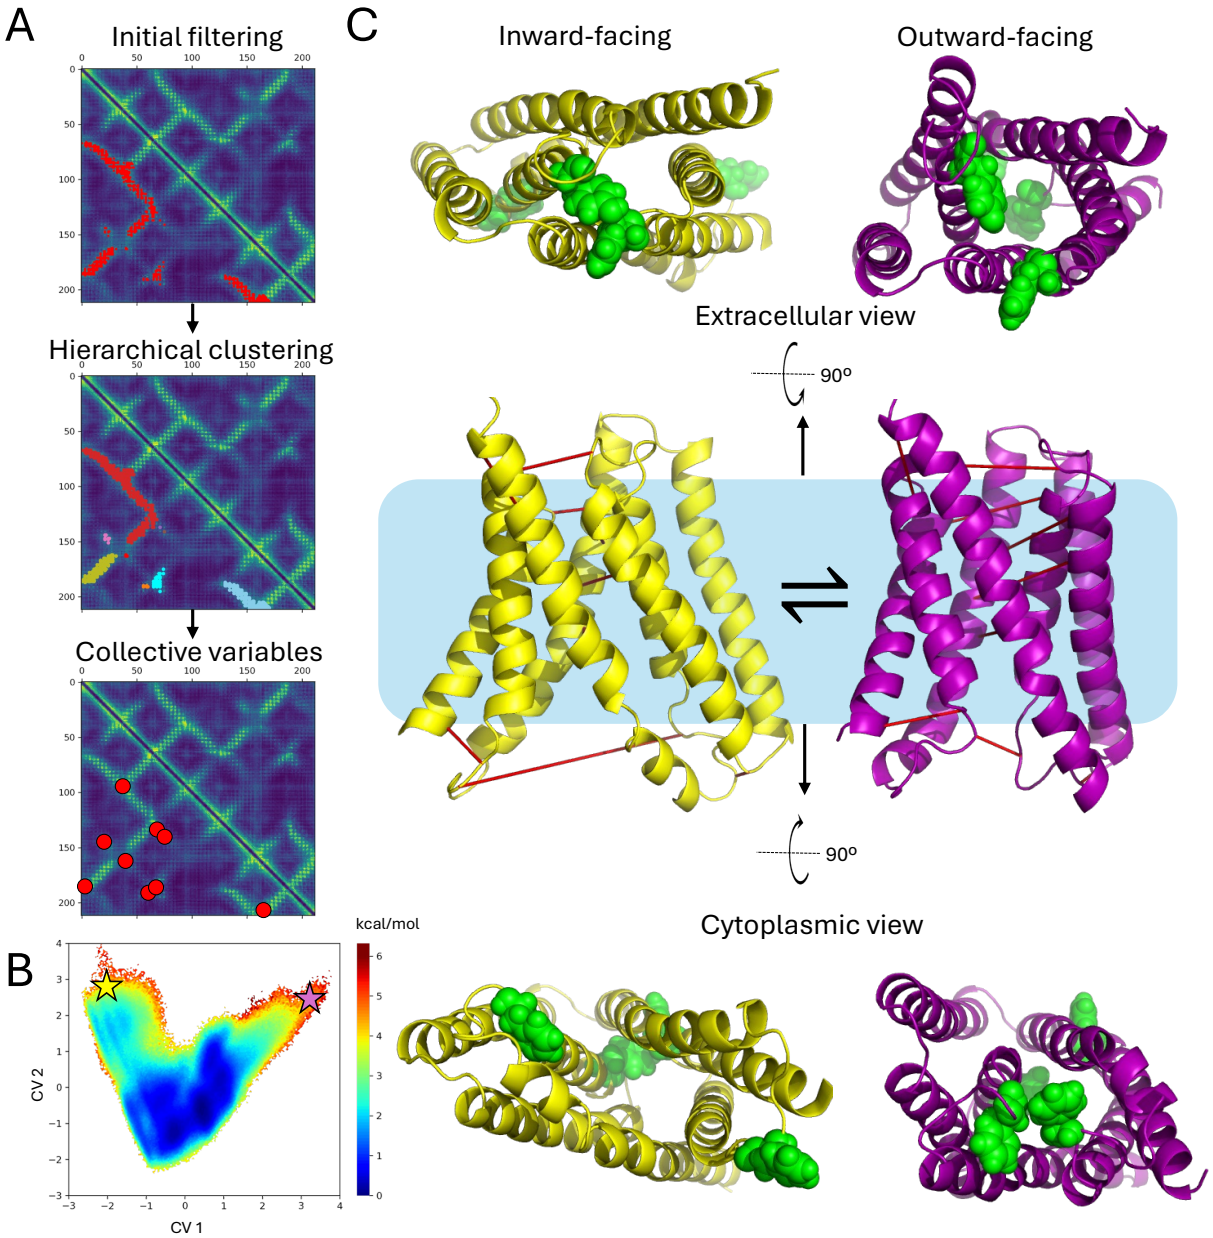

Figure 4: Automatic CV selection with ESMDynamic on the sugar transporter SWEET2b. (A) Pipeline: predicted dynamic contacts are filtered for long-range pairs, then hierarchically clustered into dynamic regions. From each cluster, the residue pair with the highest predicted probability is chosen as a CV. (B) Free energy landscape from MD trajectories using these CVs, separating inward- (yellow star) and outward-facing (purple star) states. (C) Transporter states and molecular gates: *center*, side view with CVs (red lines); *top*, extracellular view with gating residues (green) showing the extracellular gate (Y60–F191) in its closed (inward-facing) and open (outward-facing) states; *bottom*, intracellular view of the cytoplasmic gate (F42–F164) in its open (inward-facing) and closed (outward-facing) states.

ESMDynamic captures conformational switching in a *de novo* design

Recent advances in deep learning have enabled the design of multistate proteins that can dynamically transition between distinct conformations.[56–58] These systems provide a valuable benchmark for evaluating our model, as they test the model's ability to generalize features of conformational dynamics. One such example is a newly engineered variant of the N-terminal domain of troponin C, a calcium-binding protein, which was designed to adopt two alternative conformations under standard thermodynamic conditions. Upon calcium binding, the equilibrium shifts between these states.[58] This protein was deposited in structural databases after ESMDynamic training was completed, and its *de novo* nature introduces a challenging example. Unlike earlier protein design efforts that primarily focus on inducing large-scale interdomain rearrangements or rigid-body modifications,[57] this system features intradomain conformational changes, making it particularly well-suited for assessing ESMDynamic's sensitivity to subtle contact fluctuations.

For this case study, we used the I89 variant.[58] Figure 5A displays the predicted dynamic contact map, where red circles mark contacts identified as dynamic based on experimental *apo* structures (PDB IDs: 9CIG and 9CID).[58] ESMDynamic successfully captures most (94%) of the experimentally observed dynamic contacts and also predicts additional high-probability contacts in nearby regions. Notably, there is a probability-enriched segment involving the distal portion of α-helix A (residues A29–F33), a flexible loop (D34–G38), and α-helix B (Y43–L53). Although these regions appear structurally conserved between the two experimental states, local fluctuations are likely due to loop flexibility, and our findings are consistent with those shown in Figure S15 of ref. 58. Figures 5B and 5C illustrate examples of dynamic contacts in this system: the experimental structures are rendered with residue pairs M7–K68 and K44–G75 in space-filling representation. While the first contact is favored in state 1, the second is favored in state 2. Overall, ESMDynamic demonstrates excellent performance in this challenging example of a *de novo* multistate protein.

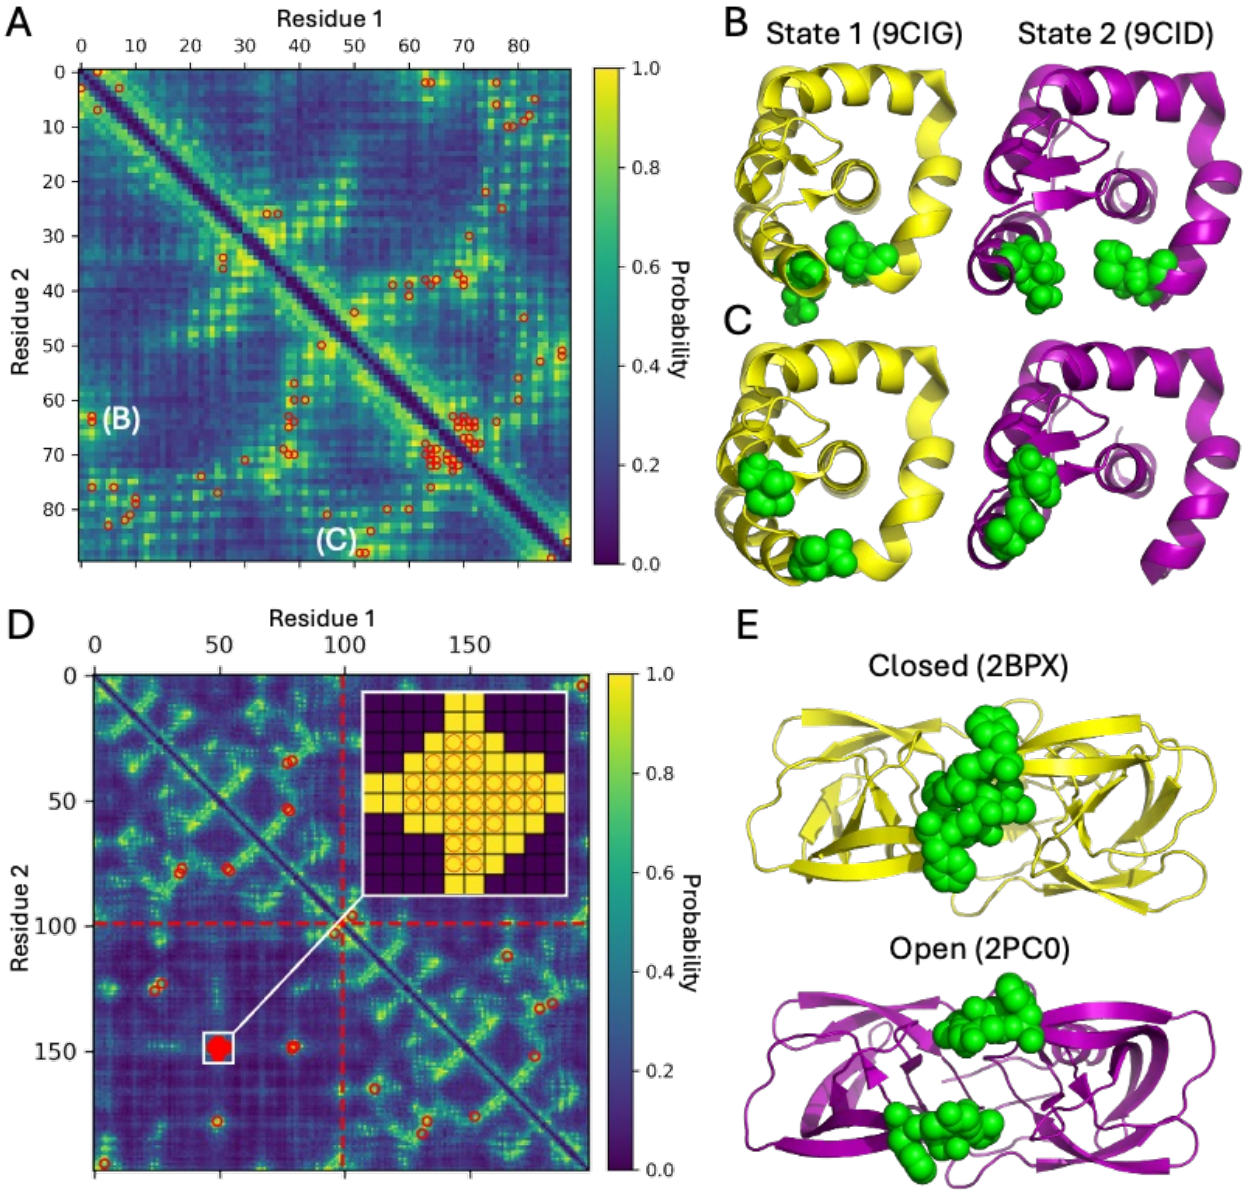

Figure 5: Additional ESMDynamic applications. (A) Predicted dynamic contact map for *de novo* multistate design "I89" from troponin C (residues indexed from A5). Red circles: dynamic contacts from experimental structures. Contacts M7–K68 and K44–G75 are labeled B and C. (B) Experimental states 1 and 2 with M7 and K68 in green. (C) Experimental states 1 and 2 with K44 and G75 in green. (D) Predicted dynamic contact map for HIV-1 protease homodimer (residues indexed from P1; chains indexed consecutively). Red circles: dynamic contacts from experimental structures; red dashed line: chain boundary. Inset shows a dynamic contact cluster at the dimer interface (I47–F53), identified by ESMDynamic as a high-probability site. (E) Experimental structures with I47–F53 in green; top: closed conformation, bottom: open.

ESMDynamic shows emergent ability to predict dynamic protein interfaces

Although ESMDynamic was not explicitly trained on multimeric proteins or protein interfaces, it inherits many features from its base model, ESMFold, which has demonstrated emergent capabilities for multimer structure prediction despite also lacking multimer-specific training.[8] We refer to this behavior as "emergent" because it arises without direct supervision on multimeric systems. In this context, it is natural to ask whether ESMDynamic can similarly predict dynamic contacts at protein–protein interfaces.

A comprehensive validation of ESMDynamic in multimeric contexts remains challenging due to limited data availability, particularly the scarcity of datasets for multimeric systems resolved in multiple conformational states or sampled extensively via MD simulations. Despite these limitations, we demonstrate a representative application of ESMDynamic to a well-characterized homodimer, the HIV-1 protease. This viral enzyme plays a critical role in HIV maturation and is one of the most effective therapeutic targets in antiretroviral therapy.[59,60] Dimer interface dynamics are central to the enzyme's function, especially the region spanning residues I47 to F53, which forms part of the binding pocket gate that must open to accommodate protease inhibitors.[61]

Given the availability of experimental structures in both closed[62] and open[63] conformations, HIV-1 protease provides a suitable example for evaluating interface dynamics. ESMDynamic predicts a high-probability dynamic contact cluster in the I47—F53 region (Figure 5D), capturing all dynamic contacts identified from the experimental structures. Additional predicted dynamic contacts within the dimer interface are also present, though some are not directly resolved in the available structures. Furthermore, ESMDynamic identifies several intrachain dynamic contacts, many of which are not observed in the two crystal structures. We attribute these discrepancies to the limited structural sampling provided by only two static conformations, whereas broader conformational variability is expected under physiological conditions.

In summary, ESMDynamic appears to exhibit emergent capabilities analogous to those of ESMFold, suggesting it may be capable of identifying dynamic contacts in multimeric systems, provided that ESMFold can produce a reliable structural model of the multichain complex.

Discussion

ESMDynamic introduces a new paradigm for understanding protein flexibility by predicting dynamic contacts directly from sequence. Unlike traditional methods that rely on computationally intensive ensemble generation or are limited to static contact inference, ESMDynamic provides a lightweight, interpretable representation of conformational dynamics in the form of residue-residue contact probabilities. This representation captures both global and local flexibility, enabling applications across structural biology and simulation-based workflows.

Our benchmarks demonstrate that ESMDynamic matches or outperforms the current state-of-the-art ensemble generation models (BioEmu, AlphaFlow, and ESMFlow) for the task of dynamic contact prediction. We attribute this improvement to the model's direct supervision on dynamic contact labels, which allows it to learn subtle sequence-derived features that govern contact variability. AlphaFlow and ESMFlow are constrained by the structure modules of their base models, limiting exploration to high-likelihood conformations. In contrast, ESMDynamic is explicitly optimized to capture both common and rare contact states, resulting in greater recall and improved balanced accuracy. Remarkably, ESMDynamic also performs comparably to BioEmu on this task at a fraction of the compute cost, despite BioEmu being trained on a much larger MD dataset.

A key strength of ESMDynamic is its generalizability across contexts. Although it was not trained on designed proteins or multimeric systems, it performed robustly in the tested examples. This is likely due to the inductive bias inherited from ESMFold, which exhibits

emergent structural generalization despite its single-chain training objective. Our case studies, including a *de novo* multistate variant of troponin C and the HIV-1 protease dimer, suggest that ESMDynamic retains this generalization capacity when predicting dynamic contacts. These results point to broader applicability in areas such as protein design, where flexibility and allosteric communication often play a critical role.

Importantly, ESMDynamic provides an interpretable bridge between sequence and dynamics that can be harnessed for downstream applications. In our example of CV selection for MSM construction, dynamic contact predictions were used to automatically identify key collective variables from a single protein sequence. This enabled high-quality kinetic modeling of a sugar transporter's functional transitions, with performance matching that of expert-curated features. The simplicity and speed of this workflow illustrate how ESMDynamic can streamline complex analyses and lower the barrier to incorporating dynamics into protein modeling pipelines.

In addition to superior accuracy, ESMDynamic offers significant practical advantages. Since the model predicts dynamic contact maps in seconds using a single forward pass, we envision that ESMDynamic will enable large-scale or iterative applications such as database annotation or reinforcement learning for sequence design. In contrast, ensemble-based generative models require hundreds of samples per input, resulting in orders-of-magnitude greater runtime.

Despite these advantages, certain limitations remain. The current training set is biased toward ordered, monomeric proteins with sufficient structural sampling, which limits the model's exposure to diverse quaternary interfaces. Moreover, ESMDynamic was not trained on intrinsically disordered proteins, and its performance is expected to degrade in cases where the structural hypothesis from ESMFold has low confidence. As more data become available, particularly from long-timescale molecular dynamics simulations, future iterations of ESMDynamic could be retrained or finetuned to better capture these challenging behaviors. Additionally, while dynamic contact probabilities offer a compact and interpretable

summary of conformational flexibility, they do not explain the underlying mechanisms of motions. Integrating ESMDynamic with physics-based computational sampling and experimental techniques will enable recovery of full conformational pathways.

In conclusion, ESMDynamic demonstrates that accurate, scalable, and generalizable prediction of dynamic residue-residue contacts is achievable directly from sequence. Its strong performance across a variety of tasks, coupled with its interpretability and efficiency, suggests broad potential for integration into both experimental and computational pipelines that require insights into protein flexibility.

Methods

Datasets

RCSB and AlphaFold DB structure ensembles. Structural ensembles were sourced from the RCSB Protein Data Bank (PDB). Both experimental structures and AlphaFold DB[37] models were utilized. MMseqs2-based groupings[6,7] were used to generate 128,660 clusters with >95% intra-cluster identity. Of these, 88,294 clusters contained at least one dynamic contact. Sequences within each cluster were aligned to the consensus sequence, and corresponding 3D coordinates were used to compute dynamic contacts (see Figure S5 for a full pipeline overview). Figures S6A,B display the cluster size distributions.

Molecular dynamics (MD) simulations. MD data were obtained from the mdCATH[38] and ATLAS[43] databases. Dynamic contacts in simulations were computed identically to those in experimental clusters, treating each simulation frame as an independent conformation. For mdCATH, only simulations at the lowest temperature (320 K) were used.

Definition of dynamic contacts. Dynamic contacts were defined as residue-residue pairs (based on C_α atoms) for which at least two structures differed across a contact threshold.

Specifically, a pair was considered dynamic if the inter-residue distance was <7.5 Å in one conformation and >8.5 Å in another, introducing a 0.5 Å margin to reduce spurious classifications from small fluctuations.

Data splits. Datasets were split into training, validation, and test sets (90%/5%/5%) using MMseqs2-based clustering.[6,7] Clusters assigned to different splits shared no more than 20% sequence identity to ensure independence. 80% minimum coverage was used (-c 0.8) for clustering. For ATLAS, we adopted the test set from AlphaFlow.[27] Full details on the data splits are provided in the GitHub repository.

Model comparisons. For model comparisons, we used the publicly available weights from `https://github.com/bjing2016/alphaflow` and `https://github.com/microsoft/bioemu`. All MSAs were generated by querying the ColabFold server[64] using the corresponding script in `https://github.com/bjing2016/alphaflow`. For the ATLAS test set, the precomputed ensembles provided by the AlphaFlow authors were employed. For BioEmu, 250 structures were sampled per target to match the number used by AlphaFlow. For the mdCATH test set, 250 structures were sampled per target using each model.

Additional simulations. MD simulations for SWEET2b were obtained from ref. 49.

Training Procedure

Overview. All training was performed on a single NVIDIA RTX 4090 GPU. ESMFold weights were kept frozen throughout (see Table S1).

Pretraining. Pretraining was performed on the experimental structure clusters. During each epoch, 10,000 samples were drawn from the training set and 1,000 from the validation set. Proteins were randomly cropped to 256 residues, with sampling probabilities weighted by protein length. Figures S6C,D show that the crop size is appropriate, since >93% of

dynamic contacts are captured. The Adam optimizer[65] was used with a learning rate of 1×10^{-4}. The focal loss function[42] was employed with parameters $\alpha = 0.99$ and $\gamma = 2$, and a batch size of 64. Training was terminated once the validation loss did not improve for 10 consecutive epochs.

Fine-tuning. Training on the mdCATH dataset followed the same procedure as pretraining, except with 1,000 training samples and 100 validation samples per epoch. To reflect the different class balance, α in the loss function was adjusted to 0.85.

Ablation experiments. For the auxiliary input ablation, transition layers were removed. For the DCM ablation, the pairwise representation from the final Evoformer block of ESMFold was used in place of the DCM output. For the prediction-layer ablation, average pooling was applied to the pairwise DCM representation prior to classification.

Validation Metrics

Balanced accuracy. To address class imbalance, we report balanced accuracy, defined as:

$$\text{Balanced Accuracy} = \frac{1}{2}\left(\frac{\text{TP}}{\text{TP}+\text{FN}} + \frac{\text{TN}}{\text{TN}+\text{FP}}\right)$$

where TP, TN, FP, and FN represent true positives, true negatives, false positives, and false negatives, respectively. This metric reflects the average of recall (sensitivity) and specificity, and treats both classes symmetrically.

Resolution metric. To evaluate the spatial precision of dynamic contact predictions, we introduce a resolution metric. For each predicted positive, the minimum L_1 distance to a true positive is computed. The distances distribution is normalized into a probability density. We observe that these distributions typically follow an exponential decay, so we fit an exponential function ($y = P \exp(-x/L)$) to obtain the precision and the characteristic decay length. A

lower decay length indicates higher spatial resolution, reflecting more accurate localization of dynamic contacts.

Data availability

Featurized datasets and model weights are available in the Illinois Data Bank at `https://doi.org/10.13012/B2IDB-3773897_V1`.

Code availability

All code used to train the model and reproduce the results of the paper is available at `https://github.com/ShuklaGroup/esmdynamic`.

A user-friendly interface to use the model is available through Google Colab Notebooks at `https://colab.research.google.com/github/ShuklaGroup/esmdynamic/blob/main/examples/esmdynamic/esmdynamic.ipynb`.

Acknowledgement

The authors acknowledge support from the National Science Foundation Early CAREER Award (NSF MCB-1845606) and the National Institutes of Health (Award No. R35GM142745). Diego E. Kleiman was supported by a fellowship from The Molecular Sciences Software Institute under NSF grant CHE-2136142.

Contributions

D.E.K., J.F., and D.S. conceived and designed the study. D.E.K., J.F., and Z.X. collected and analyzed the data. D.E.K. developed the software and machine learning model. D.E.K.

drafted the manuscript. J.F., Z.X., and D.S. contributed to manuscript revision and provided feedback. D.S. supervised the study.

References

(1) Jumper, J. *et al.* Highly accurate protein structure prediction with alphafold. *Nature* **596**, 583–589 (2021).

(2) Ahdritz, G. *et al.* Openfold: Retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization. *Nat. Methods* **21**, 1514–1524 (2024).

(3) Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature* **630**, 493–500 (2024).

(4) Chai Discovery team *et al.* Chai-1: Decoding the molecular interactions of life. Preprint at https://doi.org/10.1101/2024.10.10.615955 (2024).

(5) Wohlwend, J. *et al.* Boltz-1: Democratizing biomolecular interaction modeling. Preprint at https://doi.org/10.1101/2024.11.19.624167 (2024).

(6) Steinegger, M. & Söding, J. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat. Biotechnol.* **35**, 1026–1028 (2017).

(7) Steinegger, M. & Söding, J. Clustering huge protein sequence sets in linear time. *Nat. Commun.* **9**, 2542 (2018).

(8) Lin, Z. *et al.* Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).

(9) Hayes, T. *et al.* Simulating 500 million years of evolution with a language model.. *Science* **387**, 850–858 (2025).

(10) Hou, C., Zhao, H. & Shen, Y. Learning biophysical dynamics with protein language models. Preprint at https://doi.org/10.1101/2024.10.11.617911 (2025).

(11) Cooper, A. Thermodynamic fluctuations in protein molecules. *Proc. Natl Acad. Sci. USA* **73**, 2740–2741 (1976).

(12) Nam, K. & Wolf-Watz, M. Protein dynamics: The future is bright and complicated! *Struct. Dyn.* **10**, 1 (2023).

(13) Bacon, F. On the nature of allosteric transitions: a plausible model. *J. Mol. Biol.* **12**, 88–118 (1965).

(14) Astore, M. A., Pradhan, A. S., Thiede, E. H. & Hanson, S. M. Protein dynamics underlying allosteric regulation. *Curr. Opin. Struct. Biol.* **84**, 102768 (2024).

(15) Hyeon, C. & Onuchic, J. N. A structural perspective on the dynamics of kinesin motors. *Biophys. J.* **101**, 2749–2759 (2011).

(16) Schwartz, S. D. Protein dynamics and enzymatic catalysis. *J. Phys. Chem. B.* **127**, 2649–2660 (2023).

(17) Dutta, S., Zhao, L. & Shukla, D. Dynamic mechanism for subtype selectivity of endocannabinoids. Preprint at https://doi.org/10.1101/2024.10.25.620304 (2024).

(18) Weigle, A. T. & Shukla, D. The arabidopsis AtSWEET13 transporter discriminates sugars by selective facial and positional substrate recognition. *Commun. Biol.* **7**, 764 (2024).

(19) Nguyen, A. T., Weigle, A. T. & Shukla, D. Functional regulation of aquaporin dynamics by lipid bilayer composition. *Nat. Commun.* **15**, 1848 (2024).

(20) Zhao, C., Kleiman, D. E. & Shukla, D. Resolving binding pathways and solvation thermodynamics of plant hormone receptors. *J. Biol. Chem.* **299**, 105456 (2023).

(21) Bryan, P. N. & Orban, J. Proteins that switch folds. *Curr. Opin. Struct. Biol.* **20**, 482–488 (2010).

(22) Chakravarty, D. *et al.* Alphafold predictions of fold-switched conformations are driven by structure memorization. *Nat. Commun.* **15**, 7296 (2024).

(23) Del Alamo, D., Sala, D., Mchaourab, H. S. & Meiler, J. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *eLife* **11**, e75751 (2022).

(24) Wayment-Steele, H. K. *et al.* Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature* **625**, 832–839 (2024).

(25) Monteiro da Silva, G., Cui, J. Y., Dalgarno, D. C., Lisi, G. P. & Rubenstein, B. M. High-throughput prediction of protein conformational distributions with subsampled AlphaFold2. *Nat. Commun.* **15**, 2464 (2024).

(26) Zheng, S. *et al.* Predicting equilibrium distributions for molecular systems with deep learning. *Nat. Mach. Intell.* **6**, 558–567 (2024).

(27) Jing, B., Berger, B. & Jaakkola, T. AlphaFold meets flow matching for generating protein ensembles. Preprint at https://doi.org/10.48550/arXiv.2402.04845 (2024).

(28) Vani, B. P., Aranganathan, A., Wang, D. & Tiwary, P. AlphaFold2-RAVE: From sequence to Boltzmann ranking. *J. Chem. Theory Comput.* **19**, 4351–4354 (2023).

(29) Zhang, Y. *et al.* Distance-AF: modifying predicted protein structure models by AlphaFold2 with user-specified distance constraints. Preprint at https://doi.org/10.1101/2023.12.01.569498 (2023).

(30) Motlagh, H. N., Wrabl, J. O., Li, J. & Hilser, V. J. The ensemble nature of allostery. *Nature* **508**, 331–339 (2014).

(31) Lerner, E. *et al.* FRET-based dynamic structural biology: Challenges, perspectives and an appeal for open-science practices. *eLife* **10**, e60416 (2021).

(32) Schiemann, O. *et al.* Benchmark test and guidelines for DEER/PELDOR experiments on nitroxide-labeled biomolecules. *J. Am. Chem. Soc.* **143**, 17875–17890 (2021).

(33) Doshi, U., Holliday, M. J., Eisenmesser, E. Z. & Hamelberg, D. Dynamical network of residue–residue contacts reveals coupled allosteric effects in recognition, catalysis, and mutation. *Proc. Natl Acad. Sci. USA* **113**, 4735–4740 (2016).

(34) Mittal, S., Dutta, S. & Shukla, D. Reconciling membrane protein simulations with experimental DEER spectroscopy data. *Phys. Chem. Chem. Phys.* **25**, 6253–6262 (2023).

(35) Barducci, A., Bussi, G. & Parrinello, M. Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Phys. Rev. Lett.* **100**, 020603 (2008).

(36) Mardt, A., Pasquali, L., Wu, H. & Noé, F. Vampnets for deep learning of molecular kinetics. *Nat. Commun.* **9**, 5 (2018).

(37) Varadi, M. *et al.* AlphaFold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res.* **50**, D439–D444 (2022).

(38) Mirarchi, A., Giorgino, T. & De Fabritiis, G. mdCATH: A large-scale MD dataset for data-driven computational biophysics. *Scientific Data* **11**, 1299 (2024).

(39) Lewis, S. *et al.* Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science* **389**, eadv9817 (2025).

(40) Adhikari, B. & Cheng, J. Protein residue contacts and prediction methods. In *Data Mining Techniques for the Life Sciences*, 463–476 (Springer, 2016).

(41) Knapp, E., Fischer, S. & Parak, F. Protein dynamics from Mössbauer spectra. the temperature dependence. *J. Phys. Chem.* **86**, 5042–5047 (1982).

(42) Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988 (2017).

(43) Vander Meersche, Y., Cretin, G., Gheeraert, A., Gelly, J.-C. & Galochkina, T. ATLAS: protein flexibility description from atomistic molecular dynamics simulations. *Nucleic Acids Res.* **52**, D384–D392 (2024).

(44) Scalise, M., Pochini, L., Console, L., Losso, M. A. & Indiveri, C. The human SLC1A5 (ASCT2) amino acid transporter: from function to structure and role in cell biology. *Front. Cell Dev. Biol.* **6**, 96 (2018).

(45) Garaeva, A. A. *et al.* Cryo-em structure of the human neutral amino acid transporter ASCT2. *Nat. Struct. Mol. Biol.* **25**, 515–521 (2018).

(46) Garibsingh, R.-A. A. *et al.* Rational design of ASCT2 inhibitors using an integrated experimental-computational approach. *Proc. Natl Acad. Sci. USA* **118**, e2104093118 (2021).

(47) Garaeva, A. A., Guskov, A., Slotboom, D. J. & Paulino, C. A one-gate elevator mechanism for the human neutral amino acid transporter ASCT2. *Nat. Commun.* **10**, 3427 (2019).

(48) Yu, X. *et al.* Cryo-EM structures of the human glutamine transporter SLC1A5 (ASCT2) in the outward-facing conformation. *eLife* **8**, e48120 (2019).

(49) Selvam, B., Yu, Y.-C., Chen, L.-Q. & Shukla, D. Molecular basis of the glucose transport mechanism in plants. *ACS Cent. Sci.* **5**, 1085–1096 (2019).

(50) Singh, J. *et al.* Physiological implications of sweets in plants and their potential applications in improving source–sink relationships for enhanced yield. *Plant Biotechnol. J.* **21**, 1528–1541 (2023).

(51) Müllner, D. Modern hierarchical, agglomerative clustering algorithms. Preprint at https://doi.org/10.48550/arXiv.1109.2378 (2011).

(52) Pérez-Hernández, G., Paul, F., Giorgino, T., De Fabritiis, G. & Noé, F. Identification of slow molecular order parameters for markov model construction. *J. Chem. Phys.* **139**, 015102 (2013).

(53) Hoffmann, M. *et al.* Deeptime: a python library for machine learning dynamical models from time series data. *Mach. Learn.: Sci. Technol.* **139**, 015009 (2021).

(54) Husic, B. E. & Pande, V. S. Markov state models: From an art to a science. *J. Am. Chem. Soc.* **140**, 2386–2396 (2018).

(55) Wu, H. & Noé, F. Variational approach for learning markov processes from time series data. *J. Nonlinear Sci.* **30**, 23–66 (2020).

(56) Langan, R. A. *et al.* De novo design of bioactive protein switches. *Nature* **572**, 205–210 (2019).

(57) Praetorius, F. *et al.* Design of stimulus-responsive two-state hinge proteins. *Science* **381**, 754–760 (2023).

(58) Guo, A. B. *et al.* Deep learning–guided design of dynamic proteins. *Science* **388**, eadr7094 (2025).

(59) Nicholson, L. K. *et al.* Flexibility and function in HIV-1 protease. *Nat. Struct. Biol.* **2**, 274–280 (1995).

(60) Brik, A. & Wong, C.-H. Hiv-1 protease: mechanism and drug discovery. *Org. Biomol. Chem.* **1**, 5–14 (2003).

(61) Liu, Z. *et al.* Pulsed EPR characterization of HIV-1 protease conformational sampling and inhibitor-induced population shifts. *Phys. Chem. Chem. Phys.* **18**, 5819–5831 (2016).

(62) Munshi, S. *et al.* Rapid X-ray diffraction analysis of HIV-1 protease–inhibitor complexes: inhibitor exchange in single crystals of the bound enzyme. *Acta Crystallogr. D Biol. Crystallogr.* **54**, 1053–1060 (1998).

(63) Heaslet, H. *et al.* Conformational flexibility in the flap domains of ligand-free HIV protease. *Acta Crystallogr. D Biol. Crystallogr.* **63**, 866–875 (2007).

(64) Mirdita, M. *et al.* Colabfold: making protein folding accessible to all. *Nat. Methods* **19**, 679–682 (2022).

(65) Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. Preprint at https://doi.org/10.48550/arXiv.1412.6980 (2014).

Ethics declarations

Competing interests

The authors declare no competing interests.

Supplementary information

Supplementary Methods, Table S1, and Figures S1-6.