




EnGens: a computational framework for generation and analysis of representative protein conformational ensembles

Anja Conev , Mauricio Menegatti Rigo , Didier Devaurs , André Faustino Fonseca, Hussain Kalavadwala,

Martiela Vaz de Freitas , Cecilia Clementi , Geancarlo Zanatta , Dinler Amaral Antunes  and Lydia E. Kavraki 

Corresponding author: Dinler Amaral Antunes, Department of Biology and Biochemistry, University of Houston, 3455 Cullen Blvd, Houston 77004, TX, USA. Tel.: +1 (713) 743-9427; Fax: 713-743-3415; E-mail: antunesda@central.uh.edu; Lydia E. Kavraki, Department of Computer Science, Rice University, 6100 Main Street, Houston 77005, TX, USA. Tel.: +1 (713) 348-5737; Fax: 713-348-5930; E-mail: kavraki@rice.edu

Abstract

Proteins are dynamic macromolecules that perform vital functions in cells. A protein structure determines its function, but this structure is not static, as proteins change their conformation to achieve various functions. Understanding the conformational landscapes of proteins is essential to understand their mechanism of action. Sets of carefully chosen conformations can summarize such complex landscapes and provide better insights into protein function than single conformations. We refer to these sets as representative conformational ensembles. Recent advances in computational methods have led to an increase in the number of available structural datasets spanning conformational landscapes. However, extracting representative conformational ensembles from such datasets is not an easy task and many methods have been developed to tackle it. Our new approach, EnGens (short for ensemble generation), collects these methods into a unified framework for generating and analyzing representative protein conformational ensembles. In this work, we: (1) provide an overview of existing methods and tools for representative protein structural ensemble generation and analysis; (2) unify existing approaches in an open-source Python package, and a portable Docker image, providing interactive visualizations within a Jupyter Notebook pipeline; (3) test our pipeline on a few canonical examples from the literature. Representative ensembles produced by EnGens can be used for many downstream tasks such as protein–ligand ensemble docking, Markov state modeling of protein dynamics and analysis of the effect of single-point mutations.

Keywords: proteins, conformational ensembles, clustering, dimensionality reduction, molecular dynamics (MD), crystal structure analysis

INTRODUCTION

Proteins are the main building blocks of cells, executing a variety of functions vital to life, such as signal transduction, immune defense and DNA replication. These functions are driven by the three-dimensional arrangement (i.e. the structural conformation) of proteins [1]. However, proteins exist in a highly complex

environment and are not static entities. The following examples demonstrate that a single protein conformation is not enough to characterize important protein dynamics driving diverse functions. First, allosteric modulations, driven by mutations or drug interactions far from the protein's active site, induce conformational changes within the active site [2], which can

Anja Conev is a PhD student in the Department of Computer Science at Rice University. Her research interests include computational structural biology, applied machine learning, molecular docking and drug design.

Maurício M. Rigo is a postdoctoral researcher in the Department of Computer Science at Rice University. His research interests include bioinformatics, immunology and protein structural biology.

Didier Devaurs is an XDF research fellow in the MRC Institute of Genetics and Cancer at the University of Edinburgh. His research interest is in cross-disciplinary quantitative biomedical research.

Andre F. Fonseca is a postdoctoral researcher in the Department of Biology and Biochemistry at the University of Houston. His research interests are cancer research, immunotherapy and tumor antigen screening.

Hussain Kalavadwala is a PhD student in the Department of Biology and Biochemistry at the University of Houston. His research interests include cell biology, molecular biology and bioinformatics.

Martiela V. de Freitas is a postdoctoral researcher in the Department of Biology and Biochemistry at the University of Houston. Her research interests include bioinformatics, genomics, immunology and immunoinformatics.

Cecilia Clementi is a professor of computational physics at the Free University of Berlin. She specializes in coarse-grain and multi-scale modelling of biophysical systems to better understand cellular functions.

Geancarlo Zanatta is a professor in the Department of Biophysics at Federal University of Rio Grande do Sul (UFRGS). His research focuses on using experimental and computational structural biology to understand pharmacological molecular interactions.

Dinler A. Antunes is an assistant professor at the Department of Biology and Biochemistry of the University of Houston. His research interests span structural computational biology and immunology.

Lydia E. Kavraki is a professor of computer science and professor of bioengineering at Rice University. Her research interests span robotics, AI and biomedicine.

Received: March 11, 2023. Revised: May 23, 2023. Accepted: June 10, 2023

© The Author(s) 2023. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

modify the protein's activity [3–5]. Second, metamorphic proteins switch between drastically different folds of the same sequence, thereby performing different functions [6–9]. Finally, intrinsically disordered proteins and intrinsically disordered protein regions constitute extreme examples of highly flexible structures. They exist as highly dynamic structural ensembles failing to form a globally stable three-dimensional shape in physiological solution, thereby performing different functions [10, 11]. All these examples demonstrate the importance of comprehensively characterizing a protein conformational landscape and identifying key conformational states to understand protein function [12].

The energy landscape theory is one framework that provides an understanding of protein structure and dynamics by analyzing a protein's free energy landscape (or free energy surface – FES) as a function of a few collective variables (CVs) [13–15]. However, the exact determination of the FES for large proteins is challenging, as it requires extensive sampling of the protein's conformational space. New methods for computational protein structure prediction and simulation are emerging and there is an increased availability of protein structure datasets [16–23]. However, a full understanding of a protein's dynamics can be reached only when the dataset spans the FES sufficiently, allowing quantitative methods (such as Markov state modeling) to be applied. In this work, we do not tackle the sampling problem, as we rely on previously generated datasets. In other words, our approach focuses on structurally representative ensembles and not on statistical thermodynamic ensembles.

There is a need to rapidly extract useful information from conformational datasets without directly modeling the dynamics [24]. Subsets of conformations extracted to represent major conformational states contained within the data provide a useful conformational summary. We call such sets *representative conformational ensembles*. In this context, the term ensemble does not refer to a statistical thermodynamic ensemble commonly used in FES theory [25]. Statistical thermodynamic ensembles populate conformational landscapes and states following the laws of statistical thermodynamics, while representative conformational ensembles are smaller subsets of conformations that encompass the conformational diversity contained within larger structural datasets.

The task we address is that of extracting representative conformational ensembles from datasets of protein structures. Extracted representative ensembles can be useful for many downstream tasks such as protein–ligand ensemble docking [26], analysis of mutational effects [27] and extensive Markov state modeling of protein dynamics [28, 29]. It is important to provide sufficient analysis of the extracted ensemble to summarize important properties of each protein state (e.g. the distance between protein domains or the distance between important residues in the active site) and help derive more intuitive insights (e.g. whether a member of the ensemble represents the protein in its active or inactive form). In this work, we develop EnGens – a computational pipeline for the generation and analysis of representative protein conformational ensembles.

Sources of protein structural datasets are now diverse. The Protein Data Bank (PDB) [30, 31], first established in 1971, has experienced steady growth over the past decade. With more than 10 000 experimentally solved protein structures deposited annually, the total number of available entries to date is around 200 000. These data have allowed for new breakthroughs in the field of protein structure prediction, including machine learning techniques such as AlphaFold2 (AF2) [16], RosettaFold [17] or ESMFold [18]. The AlphaFold database contains over 200 million protein structure

predictions and the ESM Metagenomic Atlas contains 617 million predicted metagenomic structures [32]. These new developments allow researchers to collect multiple conformations of the same protein [33–36]. For example, Takei and Ishida [33] ran AF2 with multiple parameters, while Xiao et al. [34] combined sequence clustering with AF2 to generate diverse conformations for a single input sequence. Eventually, collected conformations make up datasets whose content can be summarized and analyzed with EnGens. We call such datasets 'static' to highlight the fact that the conformations they contain are independent and not derived from simulating protein dynamics.

A more extensive analysis of protein dynamics can be performed using simulations that generate so-called 'dynamic' datasets. Conformations within these datasets are not independent – they are time-ordered and form trajectories. Molecular dynamics (MD) simulations, first developed in the late 1970s, have been established as a gold standard for exploring protein dynamics [37]. Many computational packages have since been developed, including NAMD [38], GROMACS [39, 40], AMBER [41], CHARMM [42] and OpenMM [43]. MD software is becoming more accessible with python plugins and graphical user interfaces [19]. Markov State Modeling (MSM) approaches for interpreting MD simulations have recently gained popularity, although constructing MSMs can be a lengthy process requiring extensive sampling [44]. On the other hand, EnGens can be used to gain insights into the content of MD datasets without fully modeling the dynamics.

Our approach recognizes and addresses the need for a unified computational framework to help researchers summarize the vast amount of newly available structural data in an effort to understand the conformational landscape driving protein function. EnGens builds on several existing tools that have proven useful for protein structure analysis. For the computational representation of protein structure, EnGens utilizes the PDB module of BioPython [45], as well as the rich featurization module of PyEmma [46], powered by MDTraj [47]. For dimensionality reduction and clustering, EnGens provides a diverse set of algorithms implemented across deeptime [48], scikit-learn [49], UMAP [50] and SRV [51]. EnGens brings all these tools closer to the community by providing an open-source pipeline wrapped into a portable Docker image and accompanied by extensive example workflows written in Jupyter Notebooks. Additionally, EnGens implements a set of customizable interactive visualizations providing users with detailed insights into the generated conformational ensembles.

Other similar tools complement EnGens (Table S1). CoNSensX [52] generates ensembles based on available NMR data. PEnSA [53, 54] provides different metrics (Jensen-Shannon Distance, Kolmogorov-Smirnov Statistic, Overall Ensemble Similarity) for the comparison of generated ensembles. ProDy [55, 56] provides a set of algorithms for studying protein dynamics, which includes normal mode analysis. The specificity of EnGens lies in that: (1) it provides customizable PyEmma featurization for both static and dynamic datasets; (2) it contains both linear and nonlinear dimensionality reduction techniques (linear PCA [57] and TICA [58, 59]; nonlinear UMAP [50] and SRV [51]); (3) it provides different clustering methods (hierarchical, K-means and Gaussian Mixture Models); (4) it is wrapped in an accessible Docker image and includes interactive Jupyter Notebook Workflows with rich ensemble visualizations. With these unique properties, EnGens enables users to automate the generation and analysis of protein conformational ensembles. We envision EnGens as an important and useful resource for data analysis of protein structure to support researchers in the era of big data.

In the following sections, we describe methods involved in the EnGens pipeline. Note that these methods have been previously published and extensively validated [46, 50, 51, 58, 59]. Hence, validating these methods is outside of the scope of our work. Instead, we showcase the use of the full EnGens pipeline on a set of example molecules from the literature for which static or dynamic datasets are available. This involves molecules of different sizes: a large protein complex (PI3K), a peptide drug (Compstatin) and a small molecule (Nelfinavir).

METHODS

We have developed EnGens, an automated pipeline for generating and analyzing protein conformational ensembles, given a dataset of protein structures as input. Note that EnGens pipeline has two distinct use-cases: (i) processing static protein datasets (e.g. experimental structures); (ii) processing dynamic protein datasets (e.g. MD simulations).

A static structural dataset spans multiple conformations of the same protein. Static dataset could be experimentally derived and collected from the PDB or modeled computationally (using tools such as AlphaFold or Modeller). For a dataset extracted from the PDB, EnGens can be used to reveal different conformational states and extract a representative ensemble summarizing the dataset. For a dataset compiled by computationally modeling a protein and its common mutants, EnGens can describe the conformational landscape of mutants and help point out the impact of mutations.

A dynamic structural dataset is generally a trajectory derived from an MD simulation. If the simulation involves a protein with a ligand in its active site, EnGens can point out conformational changes that occur upon binding. It is important to note that for the analysis of MD-derived data much work has been done in the field of Markov state modeling [28, 60, 61]. Modeling protein dynamics is outside the scope of EnGens pipeline as its goal is only to generate and analyze the representative conformational ensemble. However, the dynamic use-case is largely inspired by insights from Markov modeling approaches. For example, one important insight is that resolving slow processes can help identify biologically relevant conformational changes. Thus, using methods related to Markov modeling helps EnGens uncover conformational states and ensembles with biological relevance.

Both static and dynamic datasets can potentially include large numbers of structures that are difficult to systematically inspect visually. To address this issue, EnGens partitions the structural dataset into clusters and extracts a representative conformation from each cluster to form a structurally diverse conformational ensemble. The EnGens pipeline is divided into four workflows that are summarized in Figure 1. Below we give an overview of these workflows and their respective goals. A detailed description of each workflow is provided in Supplementary text.

Workflow 1: extracting featured representations from raw data

The first important step in computational analysis of protein structure is finding an appropriate representation for the data. We use the term raw data to refer to the form in which protein structure is stored in databases such as the PDB. Raw data are extracted from a comprehensive experimental study or simulation, and usually contains the three-dimensional coordinates of all atoms. However, these coordinates are often redundant and noisy. To benefit from the downstream computational pipeline and avoid the effects of noise, extracting useful components

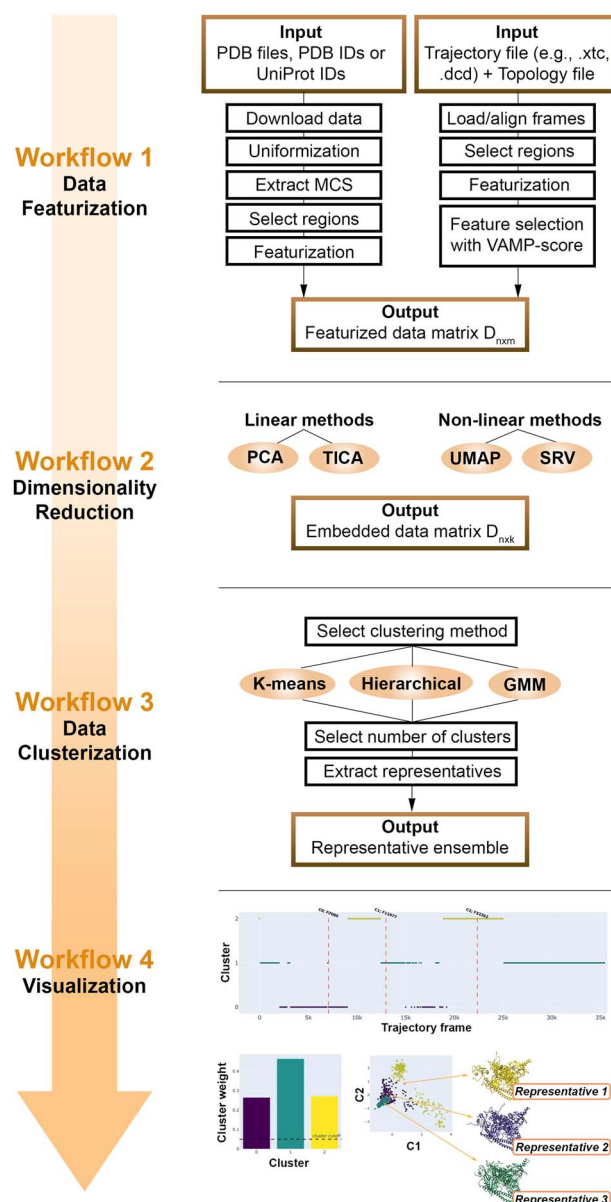


Figure 1. Overview of the EnGens methodology. Workflows are listed across the vertical arrow on the left. Individual steps are listed in the diagram on the right.

from the raw data and generating a featured representation is essential. Some features commonly used to describe the input proteins include: dihedral angles of the backbone, pairwise distances between residues and root-mean-square deviation (RMSD) distances to a reference structure. Selecting a suitable featurization remains the responsibility of the user and can be a challenging task. In Supplementary material, we provide some insights from the literature on how to wisely choose the features (Supplementary material: Workflows 1S-3 and 1D-3). The featurized representation of a protein structure is a numerical vector, which standard data science methods (such as dimensionality reduction and clustering) rely on.

Workflow 2: projecting featured representations into an embedding in low-dimensional space

Numerical vectors extracted from the first workflow often have very high dimensionality. Depending on the size of the protein and

the type of featurization, these vectors could contain thousands of elements to represent one structure. High dimensional data presents unique challenges for clustering algorithms, as metrics lose their utility in high-dimensional spaces [62]. It is thus important to embed the information into a lower dimensional space before clustering. In Workflow2, we provide implementations of four widely used algorithms for dimensionality reduction. For the static use-case, we provide two standard methods: principal components analysis (PCA) and uniform manifold approximation and projection (UMAP). For the dynamic use-case, we provide two additional methods that make use of the time-ordered nature of the data: time-lagged independent components analysis (TICA) and state-free reversible VAMPnet (SRV). TICA and SRV are not suitable for static datasets, which lack the time component that is exploited by these methods. TICA and PCA are linear methods, while UMAP and SRV can identify non linear relationships between features. The result of Workflow2 is an embedding of the data in a lower dimensional space, in which the data can be more efficiently partitioned into clusters to identify a representative ensemble.

Workflow 3: clustering embeddings and extracting the ensemble

Low-dimensional embeddings represent each conformation in the dataset. Various distance metrics can be used to calculate similarity between two conformations. This allows identifying clusters of similar datapoints. In Workflow 3, we provide implementations of three widely used clustering algorithms: hierarchical clustering [63], K-means [64] and gaussian mixture models (GMM) [65]. Hierarchical clustering provides a dendrogram of the data, allowing users to visually inspect the clusters and their relationships. The lower computational complexity of K-means makes it more suitable for large datasets. While K-means assumes a spherical data distribution, GMM can handle more complex distributions and provide a probabilistic model. Whatever the method, resulting clusters correspond to groups of structurally similar conformations. We define the hub of each cluster as the point with the most neighbors and call it a cluster representative. Eventually, we generate the conformational ensemble by extracting cluster representatives.

Workflow 4: visualizing the data and analyzing the ensemble

In the final workflow, we provide a set of customizable interactive plots to analyze the generated ensemble. Users can visualize and inspect the 2D embeddings and their clustering. The extracted representatives are highlighted and their position within the 2D embedding space can be identified. Additionally, users can visualize the 3D atomic-resolution conformations of the extracted representatives. The ensemble can be further analyzed by generating a scatterplot of interesting features (e.g. the distance between important residues or RMSD to a reference conformation) for each conformation. The same information can be summarized per cluster as a box plot. These visualizations are meant to help users interpret the ensemble (e.g. understand if the active and inactive states of a protein are represented within the ensemble).

RESULTS

The algorithms gathered under the umbrella of the EnGens pipeline have been validated in past literature [46, 50, 51, 58, 59]. The validation of these methods being therefore outside the scope of this work, in this section, we showcase the use of the full

EnGens pipeline. To this end, we have selected proteins for which structural data had been analyzed manually via often laborious processes to extract a representative conformational ensemble. We process the data entirely within the EnGens pipeline, and show that we can generate the same conformational ensemble as reported in previous studies. The examples we picked cover three systems of varying complexities. First, we process a large PI3K protein complex following both use-cases: a crystal structure dataset and an MD trajectory. Second, we apply EnGens to an MD trajectory of the peptide ligand Compstatin. Finally, we process an MD trajectory of the small drug Nelfinavir.

Class I PI3K (PI3K-I) experiments

PI3K-I is a family of lipid kinase proteins that phosphorylate a lipid found on the plasma membrane, regulating cell growth and proliferation [66]. Increased activity of PI3K-I has been associated with oncogenesis and its structural aspects have been widely studied. Members of PI3K-IA subfamily contain a regulatory (p85) and a catalytic (p110) subunit (Figure S1). Kinase activity is autoinhibited by the interaction between the nSH2 domain of the regulatory unit and the C2 domain of the catalytic unit [67, 68]. For instance, it has been shown that the nSH2 domain moves away from the catalytic unit upon contact with a phosphorylated tyrosine pY of the receptor tyrosine kinase (RTK). This movement leads to the activation of PI3K-IA [69–71]. Two recent works performed further structural analysis, one using available PI3K crystal structure data [72] and another performing and analyzing MD simulations of a mutant [73].

PI3K-IA: crystal structure dataset

We base this experiment on a study by Zhang *et al.* [72], that extracted from the PDB a dataset of 49 dimer structures corresponding to alpha, beta and delta isoforms of PI3K-IA (Table S4). While all structures are dimers (containing both catalytic and regulatory units), they differ in the portion of the regulatory unit that is crystalized, namely the nSH2, iSH2 and cSH2 domains. Ten structures were crystalized without the nSH2 domain (PI3KΔnSH2), while the rest contain the nSH2 domain (PI3K + nSH2). The analysis by Zhang *et al.* was performed using as a feature of interest the distance between the C2 and kinase domains of the PI3K catalytic unit. With a manually set threshold they divide the 49 structures into two groups: active/open (12) and inactive/closed (37). Zhang *et al.* conclude that all 10 PI3KΔnSH2 structures have nSH2 released and are active/open. Additionally, two of the PI3K+nSH2 structures have a mutation that leads to the activation. The other 37 structures are considered autoinhibited and are labeled inactive/closed.

When processing this dataset with the EnGens pipeline, our goal was to test the ability of EnGens to generate a diverse conformational ensemble that would include representative structures of the active and inactive states. We use PDB codes of the dataset as input (Table S4). EnGens extracts the maximum common substructure (MCS) for each structure (Supplementary material: Workflow 1S-2). The MCS includes the catalytic unit and the iSH2 domain of the regulatory unit (Figures S2–S4). We featurize each structure by using the pairwise distances between the centers of mass of the MCS chains. We choose the PCA option for dimensionality reduction step and K-means for clustering. Results of the analysis as provided by the EnGens dashboard are shown in Figure 2.

The dataset is clustered into five clusters. Cluster 0 contains active/open conformations of the PI3K alpha isoform (with pdb codes: 3HHM, 3HIZ and 5DXH). Cluster 1 contains eight

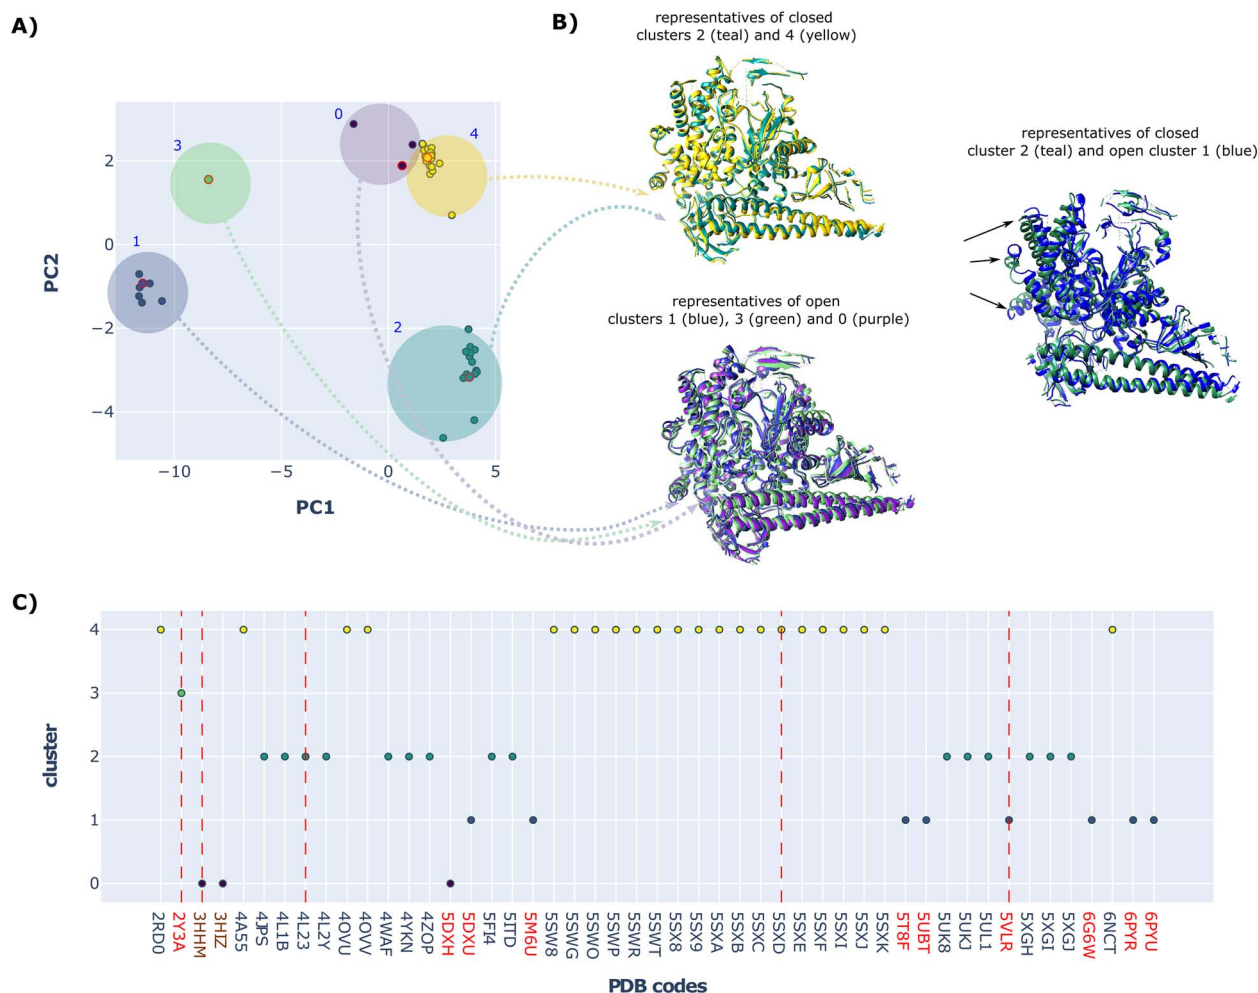


Figure 2. EnGens processing of the Zhang et al. dataset of PI3K crystal structures. A) Each point corresponds to a crystal structure. Points are colored based on the cluster they were assigned to, and clusters are indicated as large circles on the plot. Points extracted as cluster representatives are highlighted in red. The x and y axis represent the first and second principal components of these data. B) 3D structural models of the EnGens representatives: upper left - representatives of clusters 2 and 4 (inactive/closed states); bottom left - representatives of clusters 0, 1 and 3 (active/open states); right - comparison between the representative of the active state cluster 1 and the representative of the inactive state cluster 2 (the arrows point to the regions showing the biggest differences). C) PDB codes of the crystal structures are listed on the x axis. These codes are colored based on the conformational state identified by Zhang et al. (black - inactive/closed states; red - active/open states; brown - states active due to mutation). The EnGens cluster assignment is shown on the y axis. Red vertical lines indicate cluster representatives that were selected by EnGens (codes: 2Y3A, 3HHM, 4L23, 5SXD, 5VLR).

active/open conformations of the delta isoform. Cluster 3 contains a single active/open conformation of the beta isoform (2Y3A). Clusters 2 and 4 contain the remaining 37 inactive/closed conformations of the PI3K alpha isoform. The ensemble generated by EnGens contains the following representatives: 3HHM (cluster 0), 5VLR (cluster 1), 4L23 (cluster 2), 2Y3A (cluster 3), 5SXD (cluster 4). This ensemble is structurally diverse; it contains both active (3HHM, 5VLR, 2Y3A) and inactive (4L23, 5SXD) conformations. Additionally, the clusters separate the isoforms present in the dataset, namely the alpha (3HHM, 4L23, 5SXD), beta (2Y3A) and delta (5VLR) isoforms.

PI3K-IA: MD trajectory

This experiment is based on a study by Galdadas et al. [73], involving MD simulations of a PI3K-IA (with a hotspot E545K mutation leading to its increased activity), based on multiple walkers metadynamics simulations. Galdadas et al. defined two collective variables: CV1 – distance between the nSH2 domain of the regulatory unit and the helical domain of the catalytic unit; CV2 – distance to a reference state where nSH2 is detached. After inspecting the free energy surface landscape as a function of CV1

and CV2, they uncovered two energy basins: one containing a conformational ensemble corresponding to an active state with the nSH2 domain detached; the other containing two distinct conformational ensembles corresponding to an alternative activation path involving nSH2 sliding around the helical domain.

We process the MD performed by Galdadas et al. with EnGens to uncover the same conformational ensembles. To featurize the trajectory, we select: (1) the RMSD of each frame to the reference structure (first frame in the trajectory) and (2) the Cartesian coordinates of the center of mass of the helical and nSH2 domains. Next, we select SRV with a lag time of 50 to reduce the dimensionality of our input to the top 3 SRV components. We select the GMM clustering, which produces three clusters, and extract three representative conformations. The resulting EnGens dashboard is presented in Figure 3.

Clusters #0 and #1 contain conformations of the broad energy basin where the nSH2 domain is attached to the catalytic unit. Cluster #2 contains conformations in which PI3K is active and the nSH2 domain is detached. This is verified by plotting the minimum distance between residues of the helical domain (catalytic unit) and residues of the nSH2 domain (Figure 3D) for all

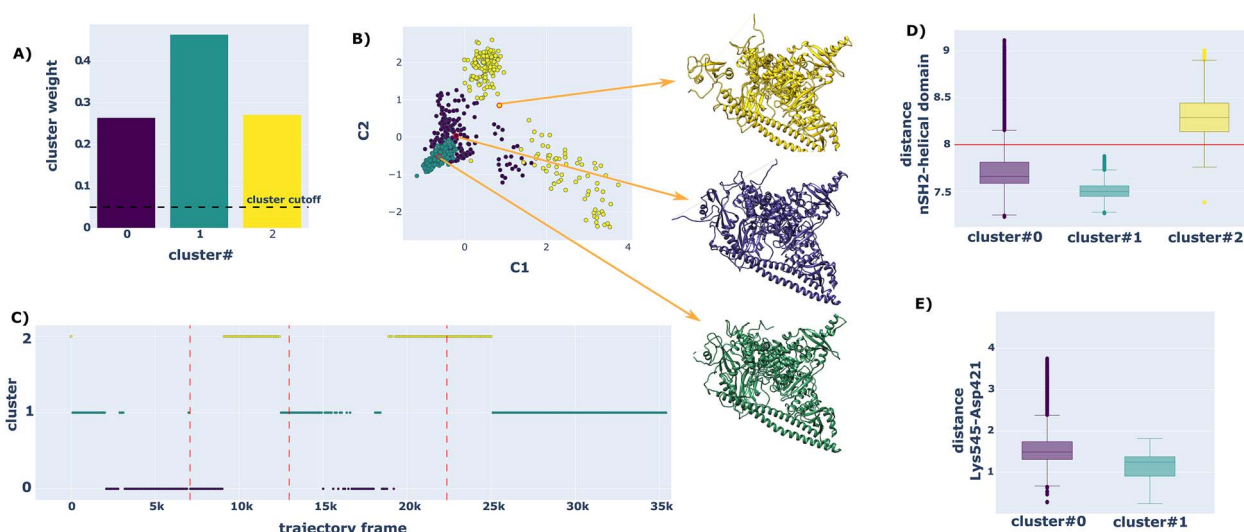


Figure 3. EnGens processing of the Galdadas et al. MD trajectory of PI3K. A) The proportion belonging to each cluster is plotted on the y-axis (cluster weight). Cluster indexes are listed on the x-axis. B) Two-dimensional embedding based on the components identified by the SRV method. Datapoints represent frames and are colored based on their respective cluster (same colors as in A). The 3D structural models of the three cluster representatives are shown on the right of this plot. C) The timeline view of the trajectory, where the x-axis lists the index of each frame, and the y-axis lists the corresponding cluster index. Vertical red lines highlight the representative frames extracted in the generated ensemble. D) The distance between nSH2 and the helical domain of PI3K is plotted on the y-axis in Å. The x-axis lists the clusters. The red horizontal line represents the threshold of 8 Å. E) The distance between Lys545 and Asp421 of the PI3K regulatory unit is plotted on the y-axis in Å. The x-axis lists the clusters.

cluster members. This distance stands out for members of cluster #2 and is higher than the 2 Å threshold. Clusters #0 and #1 contain the two conformational ensembles located in the same energy basin, as identified by the original paper. These clusters differ in the distance between the residues Lys545 of the helical domain and Asp421 of the nSH2 unit (Figure 3E). In conclusion, the three clusters identified by EnGens are consistent with the three states reported by Galdadas et al.

Compstatin experiment

Compstatin is a small, cyclic peptide that inhibits an immune surveillance mechanism associated with multiple diseases. Previous study demonstrated that compstatin analogs (i.e. biochemical variants) adopt distinct conformations that ultimately affect binding affinity and inhibitor potential [74].

We apply EnGens workflows to MD simulations obtained for two compstatin analogs, 4MeW and Cp10, chosen for their conformational heterogeneity [74]. To featurize the conformations, we select backbone torsions and carbon-alpha distances. Features are then summarized using UMAP and clustered using K-means. As a result, we retrieve representative conformations spanning several states of these analogs (Figure 4). In particular, EnGens could accurately retrieve conformational states associated with 4MeW, namely the open v -shaped state, the closed α -shaped state and three intermediate states (Figure 4A and B). These states are identified as five distinct clusters and are structurally similar to prior observations. This demonstrates again that EnGens can reproduce results obtained with distinct methodologies.

The Cp10 analog showed intriguing results (Figure 4C). EnGens produced three clusters corresponding to distinct conformations, including an intermediate state. However, Devaurs et al. had identified only two states (the opened v -shaped and closed α -shaped ones) for Cp10 [74]. We believe this discrepancy is due to limitations of this older analysis, which only relied on RMSD calculations and visual data interpretation. These new findings suggest that EnGens has high sensitivity and can capture rapid transitions between conformational states.

Nelfinavir experiment

Nelfinavir is a potent HIV-1 protease inhibitor used in adults and children. Its action mechanism involves disabling the protease from cleaving gag-pol polyprotein. However, mutations of the protease might affect the impact of Nelfinavir on patients. Using MD simulations of Nelfinavir in solution, Antunes et al. [75], inspected its conformational space and described three minimal energy Nelfinavir conformations.

We apply EnGens to these MD trajectories, using the Cartesian coordinates of all the atoms of Nelfinavir as the featurization. Then, we apply SRV for dimensionality reduction and GMM for clustering (Figure 5A). As a result, EnGens identifies nine clusters (Figure 5A and C). A cluster cutoff is used to filter out clusters containing too few points. This cutoff is similar to the one for the PI3K example (Figure 3A). In particular, using the threshold of 5%, seven out of the nine clusters are used as representative, leading to total of seven representative structures in the final ensemble (red points in Figure 5A, and red lines in Figure 5C). One cluster representative conformation coincides with the conformation described as NF-i1; other representatives are similar to the conformations described as NF-i2 and NF-i3 in the original paper (Figure 5C), considering RMSD as a measure of similarity (Table 1).

The first conformation, NF-i1, coincides with the representative of cluster 2 (RMSD = 0.413 Å). The NF-i2 structure matches clusters 1 and 6, with RMSD of 1.918 Å and 1.692 Å respectively. However, both clusters are at the end of the trajectory (Figure 5C), where they strongly overlap, indicating that EnGens slightly refined the state corresponding to the conformation presented in the original paper. The NF-i3 conformation corresponds to EnGens' cluster 3, with an RMSD of 1.336 Å.

DISCUSSION

Although protein structure prediction tools have undergone significant improvements recently, one important 'unsolved' task highlighted by the most recent CASP15 competition is modeling protein conformational ensembles. It is assumed that an




Figure 4. EnGens analysis of Compstatin analogs 4MeW and Cp10. A) UMAP visualization of the MD trajectory of the 4MeW analog. Points represent MD frames, and are colored based on their respective cluster (same colors are used in the subplot B). Cluster representatives (points highlighted in red) were selected as the closest conformation from the k-mean centroid. The 4MeW cartoon backbone for each cluster representative is presented, including three intermediate states (0, 3, 4), a single open state (cluster 1), and a single closed state (cluster 2). B) Time-oriented plot displaying the association between each MD frame (x-axis) of the 4MeW analog and the clusters. C) The Cp10 analog: cartoon backbone visualizations of the three cluster representatives, namely the intermediate (cluster 0), open (cluster 1), and closed (cluster 2) states.

Table 1. Root mean square deviation (RMSD in Å) between the cluster representatives produced by EnGens for Nelfinavir (I) and the conformations extracted from an MD trajectory (II) analyzed by Antunes *et al.* [75]. Representatives are generated by EnGens using SRV for dimensionality reduction and GMM for clustering. The lowest RMSD values in each row are reported in bold font; in the second row, two RMSD values are highlighted because they are very close to each other

II I	Cluster #0	Cluster #1	Cluster #2	Cluster #3	Cluster #4	Cluster #5	Cluster #6
NF-i1	2.150	2.879	0.413	1.436	2.837	1.475	3.682
NF-i2	3.954	1.918	3.887	3.282	3.966	4.084	1.692
NF-i3	2.304	2.701	2.467	1.336	2.572	2.292	3.563

ensemble of protein conformations will better represent the true state of a protein and will aid downstream tasks such as drug-target interaction prediction and molecular docking. Building a representative protein conformational ensemble from multiple input structures or an MD trajectory is not an easy task and many tools have been developed to tackle it. In this work, we have implemented a pipeline combining such methods, which we call EnGens.

We have evaluated the EnGens pipeline on systems of varying complexity. In each case, we recovered diverse ensembles coinciding with previously reported results. When analyzing a large protein complex such as PI3K, EnGens generated a representative ensemble containing both the active and inactive states. For the Compstatin peptide EnGens uncovered additional clusters of conformations, therefore enriching a previous study. EnGens also generated a relevant ensemble for the small drug Nelfinavir.

There are still big challenges for a pipeline of this sort. First, there are no clear guidelines on which method would perform best for a given molecular system. A number of alternative methods, each bearing its own set of hyper-parameters (Table S3), can be used to perform steps of the pipeline. We provide default values and some theoretical guidelines. For example, SRV and UMAP perform nonlinear dimensionality reduction, while TICA and PCA are linear. We thus suggest using SRV and UMAP for more complex systems where nonlinearity of features is expected. In addition, as TICA and SRV are techniques that are suitable for time-series data, they are expected to be less prone to noise resulting from fast structural fluctuations and should be suitable for the dynamic use-case. However, they can not be applied to the static use-case. Further theoretical analysis of some of these methods can be found in the literature [76]. Hyper-parameter optimization of the pipeline could be tackled with Bayesian

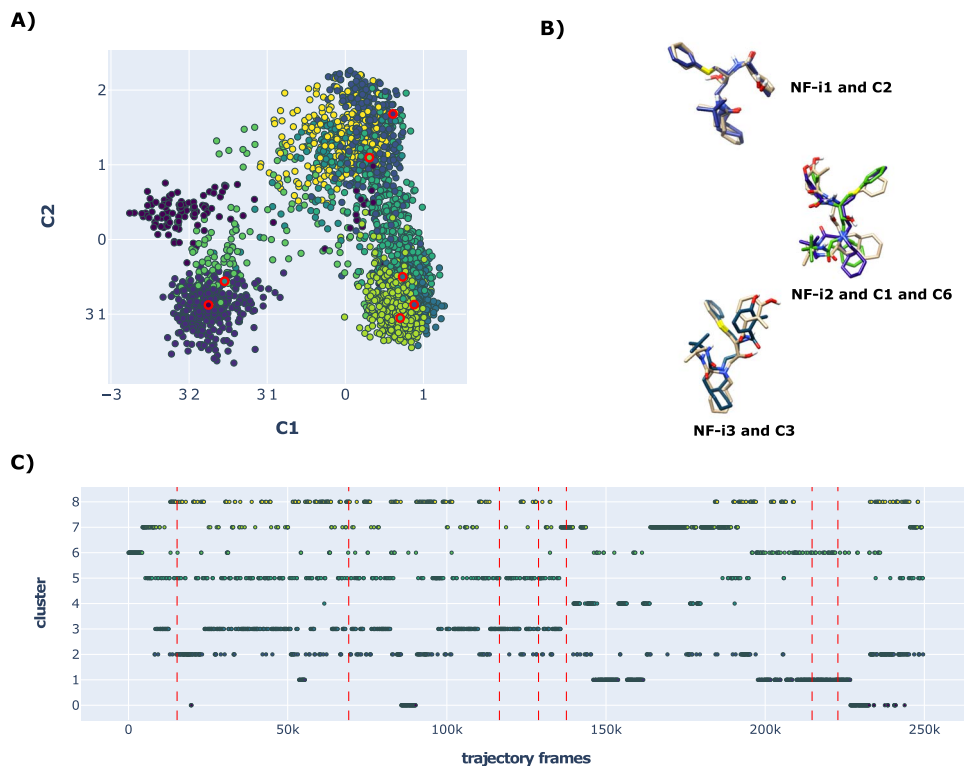


Figure 5. EnGens analysis of the Nelfinavir trajectory. A) Projection of all MD frames into a 2D space produced by SRV. Points are colored based on the cluster to which frames were assigned. B) Structural alignment between four cluster representatives and the conformations identified by [75] (NF-i1, NF-i2, NF-i3 in pale tan color). C) Timeline of the MD trajectory showing which frame (x-axis) belongs to which cluster (y-axis). Vertical red lines highlight the representative frames extracted in the generated ensemble. Points are colored based on cluster assignment using the same color scheme as in subfigure A.

optimization or other machine-learning approaches [77]. However, a wider benchmarking of these methods is necessary to evaluate the practical implications of the theory and provide good guidelines.

Second, expert knowledge of the analyzed system is still recommended for the featurization step. Some featurizations are generic, such as the pairwise residue distances that we applied to Compstatin. Others, such as the distance between the nSH2 domain and the helical domain of PI3K stem from a good understanding of the underlying system. Efforts have been made to automate this step. For the dynamic use-case new breakthroughs such as VAC (Variational Approach to Conformational dynamics) [78, 79] and VAMP (Variational Approach for Learning Markov Processes) [78] provide metrics to quantify the quality of featurization. Such metrics can be optimized using machine learning approaches to determine the most suitable featurization. However, these methods are highly dependent on the quality of the provided input MD data and are sensitive to different hyperparameters. Engineering features manually is still a widely used practice.

Third, we lack large standardized benchmarks and metrics for generating conformational ensembles. To avoid the hurdles we faced in this work, the community would greatly benefit from a public database collecting both static and dynamic datasets of protein conformations for which the representative conformational ensembles are known. Another problem is the lack of standardized metrics to compare the uncovered conformational states. Although RMSD is widely used to compare protein conformations, there are currently no equivalent standardized metrics for comparing two conformational ensembles. That is why our evaluation of EnGens is mostly qualitative and descriptive.

These challenges will become more pressing as the field moves towards big data analysis to study protein flexibility. EnGens provides easy access to existing algorithms and can serve as a platform for the rapid development of new algorithms addressing these challenges.

CONCLUSION

EnGens is a novel tool for the end-to-end processing of large protein structural datasets with the aim of generating and analyzing representative protein conformational ensembles. EnGens unifies widely used Python libraries (PyEmma, deeptime, mdtraj, UMAP, sklearn, plotly, etc.) under one Docker image and provides interactive visualizations along with extensive examples of the pipeline in Jupyter Notebook workflows. For advanced users, we provide a Python package. Our code is open source and accessible through a github repository (<https://github.com/KavrakiLab/EnGens>). We showcased how EnGens can be used to automate ensemble generation using examples from the literature. EnGens ensembles can be useful for many downstream tasks related to drug discovery such as molecular docking and drug-target interaction prediction. Additionally, EnGens can serve as a platform for further algorithmic development. Overall, we see the EnGens pipeline becoming part of many new efforts to utilize the structural data generated by novel structure prediction tools.

Key Points

- With the recent success of deep learning, some say that protein structure prediction has reached the

single-structure frontier. Conformational ensembles are said to be the future of structural biology. Datasets spanning protein conformational landscapes are becoming increasingly available. Experimental methods and generative learning models bring about static datasets, while molecular dynamics simulations provide dynamic datasets.

- EnGens (Ensemble Generator) pipeline can process both static and dynamic datasets of protein structure to extract sets of representative conformations (representative conformational ensembles) thus summarizing the complex information contained within the data. To this end, EnGens collects a long list of tools and methods previously proposed for unsupervised learning from protein structure data.
- In this work, EnGens pipeline is showcased on a set of examples from the literature including a large protein complex (PI3K), a peptide drug (Compstatin) and a small molecule (Nelfinavir). EnGens shows the ability to process these datasets faster and provide better insights than previous manual analysis.
- Representative ensembles produced by EnGens will aid downstream tasks related to drug design such as ensemble docking and drug–target interaction prediction.
- EnGens pipeline is available at <https://github.com/KavrakiLab/EnGens> as a python package, wrapped inside a Docker image and accompanied with a set of interactive Jupyter Notebooks. EnGens is accessible to researchers with little to no coding experience. For a computationally experienced audience, EnGens can serve as platform for further algorithmic development.

SUPPLEMENTARY DATA

Supplementary data are available online at <http://bib.oxfordjournals.org/>.

ACKNOWLEDGMENTS

The authors would like to thank colleagues from Kavraki Lab for many helpful discussions.

FUNDING

Work on this project by A.C. and L.E.K. has been supported in part by the National Institutes of Health NIH [U01CA258512]. Other support included: University of Edinburgh and Medical Research Council [MC_UU_00009/2 to D.D.]; Computational Cancer Biology Training Program fellowship [RP170593 to M.M.R.]; The Brazilian National Council for Scientific and Technological Development [CNPq no. 440412/2022-6 to G.Z.]; University of Houston Funds and Rice University Funds.

AUTHORS' CONTRIBUTION

A.C., M.M.R., D.D., A.F.F., M.V.F., C.C., G.Z., D.A.A. and L.E.K. conceived the experiments. A.C., A.F.F., H.K. and M.V.F. conducted the experiments. A.C., A.F.F. and M.V.F. wrote the manuscript, A.C., M.M.R., D.D., A.F.F., M.V.F., C.C., G.Z., D.A.A. and L.E.K. analyzed the results. All authors reviewed the manuscript.

DATA AVAILABILITY

The data underlying this article are available in the github repository at <https://github.com/KavrakiLab/EnGens>. The data used for the PI3K MD analysis generated by [73] are available from the EBRAINS repository at <https://search.kg.ebrains.eu/instances/Model/7f44abeb3068cdc74506bac6e72a8802>.

REFERENCES

1. Kessel A, Ben-Tal N. *Structure, Function, and Motion*, 2nd edn. Chapman and Hall/CRC, 2018.
2. Nussinov R, Tsai CJ. Allosteric in disease and in drug discovery. *Cell* 2013;**153**(2):293–305.
3. Todd AE, Orengo CA, Thornton JM. Plasticity of enzyme active sites. *Trends Biochem Sci* 2002;**27**(8):419–26.
4. Tsou CL. Active site flexibility in enzyme Catalysis. *Ann N Y Acad Sci* 1998;**864**(1):1–8.
5. Weng YZ, Chang DTH, Huang YF, et al. A study on the flexibility of enzyme active sites. *BMC Bioinformatics* 2011;**12**(1):S32.
6. Dishman AF, Volkman BF. Unfolding the mysteries of protein metamorphosis. *ACS Chem Biol* 2018;**13**(6):1438–46.
7. Dishman AF, Volkman BF. Design and discovery of metamorphic proteins. *Curr Opin Struct Biol* 2022;**74**:102380.
8. Lella M, Mahalakshmi R. Metamorphic proteins: emergence of dual protein folds from one primary sequence. *Biochemistry* 2017;**56**(24):2971–84.
9. Kim AK, Porter LL. Functional and regulatory roles of fold-switching proteins. *Structure* 2021;**29**(1):6–14.
10. Shehu A, Kavraki LE, Clementi C. Multiscale characterization of protein conformational ensembles. *Proteins* 2009;**76**(4):837–51.
11. Uversky VN. p53 Proteoforms and intrinsic disorder: an illustration of the protein structure–function continuum concept. *Int J Mol Sci* 2016;**17**(11):1874.
12. Henzler-Wildman K, Kern D. Dynamic personalities of proteins. *Nature* 2007;**450**(7172):964–72.
13. Frauenfelder H, Sligar SG, Wolynes PG. The energy landscapes and motions of proteins. *Science* 1991;**254**(5038):1598–603.
14. Kumar S, Ma B, Tsai CJ, et al. Folding and binding cascades: dynamic landscapes and population shifts. *Protein Sci* 2000;**9**(1):10–9.
15. Onuchic JN, Luthey-Schulten Z, Wolynes PG. Theory of protein folding: the energy landscape perspective. *Annu Rev Phys Chem* 1997;**48**:545–600.
16. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature* 2021;**596**(7873):583–9.
17. Baek M, DiMaio F, Anishchenko I, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science* 2021;**373**(6557):871–6.
18. Lin Z, Akin H, Rao R, et al. Evolutionary-scale prediction of atomic level protein structure with a language model. *Science* 2023;**379**:1123–30.
19. Barhaghi MS, Crawford B, Schwing G, et al. Py-MCMD: python software for performing hybrid Monte Carlo/molecular dynamics simulations with GOMC and NAMD. *J Chemical Theory Comput* 2022;**18**(8):4983–94.
20. Abraham MJ, Murtola T, Schulz R, et al. GROMACS: high performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* 2015;**1-2**:19–25.
21. Eastman P, Friedrichs MS, Chodera JD, et al. OpenMM 4: a reusable, extensible, hardware independent library for high performance molecular simulation. *J Chemical Theory Comput* 2013;**9**(1):461–9.

22. Husic BE, Charron NE, Lemm D, et al. Coarse graining molecular dynamics with graph neural networks. *J Chem Phys* 2020;**153**(19):194101.
23. Hénin J, Lelièvre T, Shirts MR, et al. Enhanced Sampling Methods for Molecular Dynamics Simulations [Article v1.0]. *Living Journal of Computational Molecular Science* 2022;**4**(1):1583.
24. Peng J, Wang W, Yq Y, et al. Clustering algorithms to analyze molecular dynamics simulation trajectories for complex chemical and biological systems. *Chin J Chemical Phys* 2018;**31**(4): 404–20.
25. Hilser VJ, García-Moreno EB, Oas TG, et al. A statistical thermodynamic model of the protein ensemble. *Chem Rev* 2006;**106**(5): 1545–58.
26. Hall-Swan S, Devaurs D, Rigo MM, et al. DINC-COVID: a web-server for ensemble docking with flexible SARS-CoV-2 proteins. *Comput Biol Med* 2021;**139**:104943.
27. Kannan A, Naganathan AN. Ensemble origins and distance-dependence of long-range mutational effects in proteins. *iScience* 2022;**25**(10):105181.
28. Abella JR, Antunes D, Jackson K, et al. Markov state modeling reveals alternative unbinding pathways for peptide-MHC complexes. *Proc Natl Acad Sci* 2020;**117**(48):30610–8.
29. Chan MC, Shukla D. Markov state modeling of membrane transport proteins. *J Struct Biol* 2021;**213**(4):107800.
30. wwPDB consortium. Protein data bank: the single global archive for 3D macromolecular structure data. *Nucleic Acids Res* 2019;**47**(D1):D520–8.
31. Burley SK, Bhikadiya C, Bi C, et al. RCSB protein data bank: powerful new tools for exploring 3D structures of biological macromolecules for basic and applied research and education in fundamental biology, biomedicine, biotechnology, bio-engineering and energy sciences. *Nucleic Acids Res* 2021;**49**(D1): D437–51.
32. Varadi M, Anyango S, Deshpande M, et al. AlphaFold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res* 2022;**50**(D1):D439–44.
33. Takei Y, Ishida T. How to select the best model from AlphaFold2 structures? [Preprint] bioRxiv 2022.04.05.487218.
34. Xiao Q, Xu M, Wang W, et al. Utilization of AlphaFold2 to predict MFS protein conformations after selective mutation. *Int J Mol Sci* 2022;**23**(13):7235.
35. Mikhaylov V, Levine AJ. Accurate modeling of peptide-MHC structures with AlphaFold. [Preprint] bioRxiv. 2023; 03.06.531396.
36. Wayment-Steele HK, Ovchinnikov S, Colwell L, et al. Prediction of multiple conformational states by combining sequence clustering with AlphaFold2. [Preprint] bioRxiv 2022; 10.17.512570.
37. Warshel A, Levitt M. Theoretical studies of enzymic reactions: dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme. *J Mol Biol* 1976;**103**(2):227–49.
38. Phillips JC, Hardy DJ, Maia JDC, et al. Scalable molecular dynamics on CPU and GPU architectures with NAMD. *J Chem Phys* 2020;**153**(4):044130.
39. Bekker H, Berendsen H, Dijkstra E, et al. Gromacs - a parallel computer for molecular-dynamics simulations: 4th international conference on computational physics (PC 92). *Phys Computing* 1993;**92**:252–6.
40. Berendsen HJC, Spoel van der D, Drunen van R. GROMACS: a message-passing parallel molecular dynamics implementation. *Comput Phys Commun* 1995;**91**(1):43–56.
41. Salomon-Ferrer R, Case DA, Walker RC. An overview of the amber biomolecular simulation package. *WIREs Comput Mol Sci* 2013;**3**(2):198–210.
42. Brooks B, Brooks C, MacKerell A, et al. CHARMM: the biomolecular simulation program. *J Comput Chem* 2009;**30**(10):1545–614.
43. Eastman P, Swails J, Chodera JD, et al. OpenMM 7: rapid development of high performance algorithms for molecular dynamics. *PLoS Comput Biol* 2017;**13**(7):e1005659.
44. Prinz JH, Wu H, Sarich M, et al. Markov models of molecular kinetics: generation and validation. *J Chem Phys* 2011;**134**(17):174105.
45. Cock PJA, Antao T, Chang JT, et al. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics* 2009;**25**(11):1422–3.
46. Scherer MK, Trendelkamp-Schroer B, Paul F, et al. PyEMMA 2: a software package for estimation, validation, and analysis of Markov models. *J Chemical Theory Comput* 2015;**11**(11):5525–42.
47. McGibbon R, Beauchamp K, Harrigan M, et al. MDTraj: a modern open library for the analysis of molecular dynamics trajectories. *Biophys J* 2015;**109**(8):1528–32.
48. Hoffmann M, Scherer M, Hempel T, et al. Deeptime: a python library for machine learning dynamical models from time series data. *Mach Learn: Sci Technol* 2021;**3**(1):015009.
49. Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: machine learning in python. *J Mach Learn Res* 2011;**12**(85): 2825–30.
50. Trozzi F, Wang X, Tao P. UMAP as a dimensionality reduction tool for molecular dynamics simulations of biomacromolecules: a comparison study. *J Phys Chem B* 2021;**125**(19):5022–34.
51. Chen W, Sidky H, Ferguson AL. Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets. *J Chem Phys* 2019;**150**(21):214114.
52. Ángyán AF, Szappanos B, Perczel A, et al. CoNSEnsX: an ensemble view of protein structures and NMR-derived experimental data. *BMC Struct Biol* 2010;**10**(1):39.
53. Vögele M, Thomson NJ, Truong ST, et al. Systematic Analysis of Biomolecular Conformational Ensembles with PENZA, 2022; arXiv:2212.02714 [q-bio.BM].
54. Vögele M, Thomson NJ, McAvity J, et al. drorlab/pensa: PENZA 0.2.8, 2022. Zenodo 4362136.
55. Bakan A, Meireles LM, Bahar I. ProDy: protein dynamics inferred from theory and experiments. *Bioinformatics* 2011;**27**(11):1575–7.
56. Zhang S, Krieger JM, Zhang Y, et al. ProDy 2.0: increased scale and scope after 10 years of protein dynamics modelling with python. *Bioinformatics* 2021;**37**(20):3657–9.
57. Pearson K. LIII. On lines and planes of closest fit to systems of points in space. *Lond Edinb Dublin Philos Mag J Sci* 1901;**2**(11): 559–72.
58. Pérez-Hernández G, Paul F, Giorgino T, et al. Identification of slow molecular order parameters for Markov model construction. *J Chem Phys* 2013;**139**(1):015102.
59. Schwantes CR, Pande VS. Modeling molecular kinetics with tICA and the kernel trick. *J Chem Theory Comput* 2015;**11**(2):600–8.
60. Husic BE, Pande VS. Markov state models: from an art to a science. *J Am Chem Soc* 2018;**140**(7):2386–96.
61. Bernetti M, Masetti M, Recanatini M, et al. An integrated Markov state model and path Metadynamics approach to characterize drug binding processes. *J Chem Theory Comput* 2019;**15**(10): 5689–702.
62. Novinskaya A, Devaurs D, Moll M, et al. Improving protein conformational sampling by using guiding projections. In: 2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), Washington, DC, USA, 2015, pp. 1272–9.
63. Murtagh F, Contreras P. Algorithms for hierarchical clustering: an overview. *WIREs Data Mining and Knowledge Discovery* 2012;**2**(1):86–97.

64. Hartigan JA, Wong MA. Algorithm AS 136: a K-means clustering algorithm. *J R Stat Soc Ser C Appl Stat* 1979;**28**(1):100–8.
65. Lindsay B, McLachlan GL, Basford KE, et al. Mixture models: inference and applications to clustering. *J Am Stat Assoc* 1989; **84**:337.
66. Martini M, De Santis MC, Braccini L, et al. PI3K/AKT signaling pathway and cancer: an updated review. *Ann Med* 2014;**46**(6): 372–83.
67. Yu J, Wjasow C, Backer JM. Regulation of the p85/p110 α phosphatidylinositol 3-kinase: distinct roles for the N-terminal and C-terminal SH2 domains. *J Biol Chem* 1998;**273**(46):30199–203.
68. Miller MS, Schmidt-Kittler O, Bolduc DM, et al. Structural basis of nSH2 regulation and lipid binding in PI3K α . *Oncotarget* 2014;**5**(14): 5198–208.
69. Buckles TC, Ziemba BP, Masson GR, et al. Single-molecule study reveals how receptor and Ras synergistically activate PI3K α and PIP3 Signaling. *Biophys J* 2017;**113**(11):2396–405.
70. Nolte RT, Eck MJ, Schlessinger J, et al. Crystal structure of the PI 3-kinase p85 amino-terminal SH2 domain and its phosphopeptide complexes. *Nat Struct Biol* 1996;**3**(4):364–74.
71. Vadas O, Burke JE, Zhang X, et al. Structural basis for activation and inhibition of class I phosphoinositide 3-kinases. *Sci Signal* 2011;**4**(195):re2.
72. Zhang M, Jang H, Nussinov R. Structural features that distinguish inactive and active PI3K lipid kinases. *J Mol Biol* 2020;**432**(22):5849–59.
73. Galdadas I, Gervasio FL, Cournia Z. Unravelling the effect of the E545K mutation on PI3K α kinase. *Chem Sci* 2020;**11**(13):3511–5.
74. Devaurs D, Antunes DA, Kaviraki LE. Computational analysis of complement inhibitor compstatin using molecular dynamics. *J Mol Model* 2020;**26**(9):231.
75. Antunes DA, Rigo MM, Sinigaglia M, et al. New insights into the in silico prediction of HIV protease resistance to nelfinavir. *PLoS One* 2014;**9**(1):e87520.
76. Glielmo A, Husic BE, Rodriguez A, et al. Unsupervised learning methods for molecular simulation data. *Chem Rev* 2021;**121**(16): 9722–58.
77. Lee Y, Chamzas C, Kaviraki LE. Adaptive experience sampling for motion planning using the generator-critic framework. *IEEE Robot Autom Lett* 2022;**7**(4):9437–44.
78. Wu H, Noé F. Variational approach for learning Markov processes from time series data. *J Nonlinear Sci* 2020;**30**(1):23–66.
79. Lorpaiboon C, Thiede EH, Webber RJ, et al. Integrated variational approach to conformational dynamics: a robust strategy for identifying Eigenfunctions of dynamical operators. *J Phys Chem B* 2020;**124**(42):9354–64.