Information theoretics for the machine learning detection of functionally conserved and coordinated protein motions

Gregory A. Babbitt[1]

[1]Thomas H. Gosnell School of Life Sciences, Rochester Institute of Technology, Rochester NY, USA 14623

Contact Information –Gregory A. Babbitt (gabsbi@rit.edu)

Short title – machine learning and molecular dynamics

Abstract

The application of machine learning classification to the molecular dynamics of the functional states of protein allows for application of an information theoretic framework similar to traditional bioinformatics. The functional states of proteins involving binding interactions with partners comprised of protein, DNA or small molecules can first be defined in a binary fashion (i.e. bound vs unbound), subsequently simulated in molecular dynamics software, and then employed as a comparative training set for a binary machine learning classifier capable of discerning the complex dynamical consequences of binding interaction. This learner can subsequently be deployed on new simulations of the functionally bound state to validate its ability to recognize the molecular motions that are supporting binding function. Regions of proteins with functionally conserved dynamics will induce significant local correlations in learning performance across independent validation runs. Through case studies of Rbp subunit 4/7 interaction in RNA Pol II and DNA-protein interactions of TATA binding protein, we demonstrate this method of detecting functionally conserved protein dynamics. We also demonstrate how Shannon information, relative entropy and mutual information can be applied to these binary classification states of dynamic simulations in order to compare dynamics and identify coordinated motions involved in dynamic interactions across sites.

Keywords – molecular dynamics, machine learning, functional conservation, molecular evolution, information theory, protein allostery

Introduction

All variety of functions of living organisms are defined by dynamic processes that proceed over a vast multitude of size and time scales. At the systems level, a biological organism can be defined by a dynamic network of genetic and environmental interactions. Within this network are dynamic signaling pathways that form feedbacks to create temporary homeostatic stable states. In turn, each of these pathways is comprised of many dynamic molecular encounters between protein, nucleic acid and small molecule structures that bind to one another via weak and strong chemical bonding interactions. And finally, each one of these biomolecules is itself a rapidly dynamic conformation of soft matter that is thermally responsive to continuous alteration of its local solvated environment. Therefore, the study of biological function within living systems would seem to require a thorough analysis of the time series of molecular structural conformation (i.e. molecular dynamics), ultimately providing the biologist a dynamic perspective on functional molecular phenotypes. The process of organic evolution also works across these scales by first slowly accumulating chemical mutations that affect the rapid dynamics of molecular processes, and then selecting the resulting self-organized complex phenotypes that are more fit and fixing them within populations. Therefore, we should expect that when the molecular dynamics of a given protein is functionally adaptive in some way, it should produce a tendency to move in a non-random and specifically directed way, much like a machine, that is readily distinguishable from thermal noise imparted by its solvent environment. It should also preserve this characteristic motion over deep evolutionary time. Therefore, when protein sequences are functionally conserved over evolutionary time scales (i.e. millions of years), we should expect to also see molecular dynamics that are conserved over very short functional time scales (i.e. nanoseconds) as well.

Unfortunately, this necessary molecular dynamic perspective on biological function has been relegated to only a few computationally heavy fields in biology (i.e. biophysics, computational biochemistry and systems biology), and is still largely ignored by a mainstream biology that maintains a methodological focus on the comparative interpretation of static representations of biomolecular sequences, sequence annotations and protein structures housed within ever-growing online databases. While molecular biology has often been generally defined as the subset of biochemistry in which information can be substantiated and manipulated (1–4), the concept of information in biology has often been historically disconnected from its mathematical relationship with randomness in thermodynamics (5) and further muddied with mathematically undefined terminology like information 'flow', 'transfer' or 'processing' (6) when referring to the molecular templating involved in the replication of DNA, the transcription of DNA to RNA and the translation of RNA to protein. Fortunately one exception has been the common application of information theory to multiple alignments of binding sequences (7,8). We feel there is room to expand the mathematical formalism that already exists around this concept of information content embodied in functionally conserved biomolecular sequences to the concept of information embodied within functionally conserved motions in biomolecular dynamic simulations as well. We would go further to argue that analyses that largely ignore dynamics, rooted only upon a sequence representation of molecular biology will continue to potentially obscure a very large component of latent functional variability that connects the genotype to the phenotype through the scaling of self-organized dynamic physicochemical processes that build the complex body forms and functions that biologists define as life (9).

While all-atom and coarse-grained molecular dynamic studies are often used to address functional chemistry and thermodynamics of many biomolecular systems, the lack of widespread adoption of dynamic simulation by the mainstream disciplines of biology can perhaps be linked to both a lack of standard statistical analytical methods for comparing molecular dynamic simulations to each other as well as this lack of a connection between information theory and molecular dynamics as they apply to biological systems. Comparative methods are central to laboratory experimentation in molecular biology, and the information-theoretical perspective applied to DNA sequences continues to enable major advances in the fields of genomics and molecular evolution that have developed in parallel with major advances in modern sequencing technologies. Concurrent technological advances in computer hardware have also allowed larger, longer and more highly parameterized (i.e. accurate) molecular dynamic simulations (10,11), and now greatly enhance our ability to investigate the link between molecular dynamics and evolutionarily-conserved function (9,12–14). However, very few methods and software currently exist allowing the exploration of function and evolution through molecular dynamics simulation (15,16). Here, we propose a generalizable method of employing machine learning-based classification of functional dynamics under the same information theoretic framework that has proven so useful for the comparative analysis of multiple-aligned protein and DNA sequences (7,8). Our method allows for the detection of functionally conserved dynamic motions using comparative molecular dynamic simulations of proteins in different functional states, without the need to resort to additional sequence-based molecular evolutionary analyses (17). While sequence-based tests of selection can readily identify adaptively altered or functionally constrained sequences in genomes, unless they are complimented with functional studies, structural analysis, or physical simulations they generally fail to identify what exactly are the specific protein functions are the target of natural selection. Comparative molecular dynamic statistical methods inherently capture specific protein function that is defined by non-random machine-like motion in the protein system. When paired with information theoretics applied to molecular genetic variants, comparative molecular dynamics can potentially supply a more complete functional evolutionary picture of variation within protein space. Information theoretics applied to binary machine learning classifications of molecular dynamic simulations can also allow for identification of coordinated motions across sites caused by structural symmetries and/or allosteric interactions. Here, we first present a general computational protocol and an information theoretical framework for analyzing functionally conserved and coordinated motions of proteins by applying machine learning to molecular dynamic simulation. We then offer two case studies employing this new approach by analyzing the binding interaction of the Rbp4/7 subunits of RNA Pol II; a well-studied protein interaction involving Rbp4 N and C terminal domains, well conserved at the sequence level, and connected by a non-conserved linker region with strikingly different lengths in humans and yeast. We find a strong similarity between functionally conserved dynamic regions of Rbp4 discovered by our method, and functionally validated regions of sequence conservation reported by others (18–23). We supplement this with an analysis of functional DNA interaction of the human TATA binding protein, highlighting a pervasive dynamic interactions that share mutual information across all sites of the protein when in its DNA bound state.

Methods

A practical simulation protocol for identifying local regions of proteins with functionally conserved dynamics

We outline a general protocol for combining molecular dynamics simulation software with machine learning packages for general programming languages with a goal of detecting functionally conserved protein dynamics related to functional binding interactions. Note: this is also the protocol implemented in our most recent software release (15). Throughout the rest of the article we abbreviate this protocol as FCDA (Functionally Conserved Dynamic Analysis)

A) Utilizing one of several popular molecular dynamic (MD) simulation software suites (e.g. Amber, NaMD, Charm or OpenMM), generate two large ensembles of MD production runs for training one or more machine learning classifiers. Each ensemble should represent a functional state of the molecular system under investigation (e.g. bound vs unbound state of a protein that interacts with a nucleic acid, another protein, or a functional small molecule ligand such as ATP). See Figure 1A. To avoid pseudo-replication and to mitigate the potential driving effects of the initial conditions of the simulation, each MD production run should start at a different point in time.
B) Extract one or more features of the dynamics for training using vector trajectory analysis software such as ptraj or cpptraj (24). Features for analysis might include both directionless features representing local vibration/flexibility (e.g. root mean square atom fluctuation) and/or directional features of vectors representing dynamic changes in protein shape as the simulations progress (e.g. XYZ vector distances to a central feature of the protein or solvent box). These data should be extracted over many time slices to compose a univariate or multivariate feature vector for the chosen machine learning classification model. See equation 4 in (15).
C) At single amino acid resolution, train one or more machine learning classifiers implemented in R or Python on these features (i.e. the feature vector) and calculate their overall learning performance on new MD production runs that represent one of the two training states of the simulation. NOTE: learning performance calculations and information content are inter-related as they both rely upon the probability of classification into one group versus the other. In simple terms, learning performance of algorithms will increase when there is information about the motions of functional states to be learned (Figure 1B). At this point one can validate the effectiveness of the training by conducting a new MD validation run on the functional bound state of the protein and calculate local frequencies of correct classification at each residue across the whole protein.
D) To identify local regions with significant functionally conserved dynamics, calculate the local correlation of learning performance on two or more of the MD validations conducted in step C using a sliding window (Figure 1B). The statistical significance of this local correlation can be used to call significance to the mapping of functionally conserved dynamics. The logic here is simple. If dynamic states of local regions are occurring randomly due only to thermal noise, then learning performance will not locally correlate across MD validation runs. However, if the local amino acid residues are inducing a non-random motion that is detectable by the learner, then a local correlation in the learning performance profile will be induced. If multiple machine learners are used a canonical correlation analysis and associated significance via Wilk's lambda can be used within the sliding window (25)(Figure 1C). Multiple test correction for false discovery rate should also be employed to adjust p-values in accordance for the number of sites on the protein (26). Thus local regions with functionally conserved dynamics can be mapped as either sequence annotations or color mapped protein structures.

A common information theoretic framework for functional binding or protein sequences and the dynamics of their interactions

The concept of information in relation to entropy is derived from Shannon's work 'The mathematical theory of information' (5).

Shannon Information content (I) is defined as

$$I = I_{max} - H \tag{1}$$

Where I_{max} is the maximum information content allowed and H is Shannon entropy observed and is defined as

$$H = -\sum_{i=1}^{M} p_i * \log_2(p_i) \tag{2}$$

Where *p* is a set of discrete probabilities of *i* symbols from a set of possible discrete states (e.g. an alphabet) of size *M*, used to communicate a message over a discrete channel with noise (5).

When applied to a multiple sequence alignment of binding sites of a protein coding DNA sequence, where the frequencies of nucleobases are taken as estimates these discrete probabilities, the information content at a given site in the DNA sequence follows (7)

$$I = log_2\,(4) - H = 2 - H \tag{3}$$

where

$$H = -\sum_{b=A,T,C,G}^{4} f(b,L) * \log_2 f(b,L) \tag{4}$$

where *f(b,L)* is frequency of a given nucleobase (b) at position (L) in the best alignment or

$$f(b,L) = \left(\sum_{i=1}^{N} b\right)/N \tag{5}$$

and where N is the number of aligned sequences.

This is the method that underpins the height of letter base symbols on the vertical Y axis in the well-known sequence logos plot (8). And this approach can be applied to characters representing proteins by replacing the 4 possible nucleobase character states with the 20 possible amino acid character states. To obtain information for a given sequence rather than a given position within a sequence, one can average the information gained (i.e. entropy lost) by the best alignment over the whole sequence, assuming the frequencies of the nucleobases occurring at sites are independent. Therefore, the information content represents the total entropy lost as the best alignment of the

sequences representing the functional site is obtained. Regions of distantly related genomes with high functional conservation will maintain high information content when aligned, as these regions are where purifying selection has not allowed random nucleobase changes to accumulate over deep evolutionary time scales.

Now if information can be imagined within the context of biomolecular sequences that represent a binding interaction, we need to now ask whether we can imagine something similar in the context of molecular motion, or dynamics, over shorter time scales as well. In many ways, the notion of a molecular form changing over time is much more amenable to Shannon's original mathematical conceptualization of communication signal that also plays out over a relatively short time and must be distinguished from random noise. In our case, the motions of protein can be also defined in terms of 'signal' represented by the repeated micro-machine like motions that directly relate to a protein binding function, and 'noise' represented by random thermal motions caused by the continuous random collision events of solvent molecules hitting the protein. However, to apply the mathematical formalism of information theory within the context of functional protein dynamics, we obviously would need some previous knowledge of the function of the protein structure under examination. Protein function is most often defined by its binding interactions with nucleic acids, signaling molecules or other proteins, so if we could watch the molecular dynamics of a protein in two functional states (e.g. bound and unbound to a particular binding partner in a signal transduction pathway), and later apply machine learning to identify these functional dynamic states when the protein dynamics is restarted from different initial or novel conditions, then we might define the local information content in a functional dynamic context for the motion of a given residue on the protein as

$$I = log_2\,(2) - H = 1 - H \tag{6}$$

where

$$H = -\sum_{c=0,1}^{2} f(c,L) * \log_2 f(c,L) \tag{7}$$

and where f(c,L) is the frequency of binary classification (c) assigned by a machine learner properly trained to recognize the dynamics of the functional state (e.g. 1 = bound, 0 = unbound) at residue location (L) taken over (T) a set of discrete time slice intervals through the total length of the molecular dynamic simulation. Therefore, the frequency of classification returned upon the deployment of a learner on a molecular dynamic simulation is

$$f(c,L) = \left(\sum_{t=50ps}^{T} c\right)/T \tag{8}$$

where t is the length of any given time slice interval (50 picoseconds by default) and T is the total number of time slice intervals in the simulation. An average frequency of 0.5 (i.e. the mid-point between 0 and 1), is therefore indicative if either a poorly trained learner or a lack of local functional dynamic information to be learned from a given molecular dynamic simulation.

Modern supervised machine learning algorithms are designed to derive classification states when trained on adequate amounts of real data, and molecular dynamics simulations are capable of

producing large amounts of data on simulated motion due to its necessity to take very small step sizes. Hypothetically, if one applies any proper machine learning method of classification, trained on the simulations of functional dynamic states of interest, one can determine local information regarding functionally conserved protein dynamics without resorting to traditional sequence-based bioinformatics. Note that because the training classes are completely specified, the term $f(c, L)$ is equivalent to a traditional measure of machine learning performance where the sum of true positive and true negative classifications divided by the sum of all total classifications (i.e. true positive/negative and false positive/negative). It is important to also note, that learners should be applied independently to the dynamics of individual amino acids as it is at this resolution of single amino acid replacement, that purifying selection will act over deep time.

Mutual information between local learning performance profiles on independent identical simulations can identify regions of proteins with significantly functionally conserved dynamics

Assuming a machine learner is adequately trained on each amino acid to successfully discern local differences in functional dynamics states (i.e. when the protein is bound to a functional partner or not), then we can imagine running new simulation repeatedly in one of the two functional states (i.e. bound or unbound) and assessing whether local peaks in learning performance are correlated in their position along the protein backbone as opposed to just random in position. If local molecular dynamic motions of atoms are not functional as defined by the learner, there should be no local positional peaks in the learning performance or f(c,L) (Figure 1C). In this situation the average frequency of classification should be 0.5. Random thermal differences and random sampling events in new identically prepared simulations will not allow independent learning profiles collected from the same sort of simulation to be correlated. But if non-random functional motions are present on the protein backbone and the local learners can classify them according to functional states upon which the were trained, then only sampling errors will create local differences in learning performance causing correlated profiles in subsequent validation runs (Figure 1D). When multiple machine learning methods are used to reduce any artifacts caused by the sensitivity of a particular method of machine learning classification, then a multivariate or canonical correlation analysis (CCA) is required. High local R values from the CCA, generating Wilk's lambda with a significant p-value can be used to call or map regions with significantly high local non-random dynamics associated with function (in most cases a specific binding interaction). Given two vectors representing classification states (0,1) called from *n* different machine learning methods where X = $\{x_1, x_2, \ldots x_n\}$ and Y = $\{y_1, y_2, \ldots y_n\}$ and new variables U and V defined via linear combinations of X and Y, the CCA seeks to find the highest canonical correlation (CC) by maximizing

$$CC = corr(U, V) = \frac{cov(U, V)}{\sqrt{var(U)} * \sqrt{var(V)}} \tag{9}$$

One can see great similarity of eqn 9 to mutual information in eqn. 16. Thus the CCA essentially defines the local mutual behavior of homologous amino acid residue dynamics between two MD validation runs. From a functional evolutionary perspective, this mutual dynamic behavior identified by the learner can only arise from non-random or molecular machine-like motion that is recurrent in the time scale of the simulation. It can be implied that the ability of a protein to conduct this motion must be also conserved over deep evolutionary time scales as well.

In summary, the machine learning performance of a binary classification algorithm, well trained on the functionally bound and unbound dynamic states of a given protein, can be utilized to identify functionally conserved protein dynamics in both time and space. And, because the output of this analysis is binary (i.e. 0 and 1), information theory can be subsequently applied to address many interesting problems that involve comparative differences in dynamics across sites or states, as well as the coordination of dynamic states across distances define by the protein's structure.

One might argue that the original sequence-based alignment approach best captures the evolutionarily conserved function because the multiple sequence alignment in of itself represents molecular changes accrued over deep time scales. However, in traditional molecular evolutionary studies it is often quite challenging to relate evolutionarily-conserved regions of the genome with specific molecular functions (27,28). A machine learning enabled molecular dynamics approach to the problem can potentially detect non-random motions within conserved regions and when applied to molecular dynamic simulations of mutated proteins, it can also provide evolutionary perspective that directly speaks to the specific functional context of the molecular system under investigation.

Other applications of information theory to molecular dynamics

The information theoretic concept of relative entropy (also called discrimination information, information gain, or Kullback-Leibler divergence) (29) can be very useful when applied to comparative questions about molecular dynamics in different functional or genetic states. In essence, the relative entropy can be conceptualized as a distance measure between two distributions of interest that is sensitive to the difference in mean, spread and shape between distributions. It can be applied to direct pairwise comparisons of the distributions of the dynamic motions themselves, or alternatively it can be applied to comparisons of local correlations in machine learning performance (i.e. information content) in validation MD runs in order to assess the impacts of molecular variation on functionally conserved dynamics. This molecular variation can represent different genetic mutations, chemical epigenetic alterations variations in small molecule binding interaction (i.e. any molecular variation of interest in the simulation of a solvated protein system).

A) Relative entropy between local protein motions related to functional dynamic states

The Kullback-Leibler divergence or relative entropy between two discrete probability distributions P and Q on probability space X is given by

$$D_{KL} = \sum_{x \in X} P(x) \log\left(\frac{P(x)}{Q(x)}\right) \tag{10}$$

The symmetric form is the Jenson-Shannon divergence which defines the similarity between the two distributions

$$D_{JS} = \frac{1}{2} D_{KL}(P||M) + \frac{1}{2} D_{KL}(Q||M) \tag{11}$$

where M = ½(P+Q)

The relative entropy (RE) or similarity between the root mean square fluctuation (rmsf) of two homologous atoms moving in two molecular dynamic simulations representing a functional binding interaction (i.e. where 0 = unbound state and 1 = bound state) can similarly be described by

$$RE_{functional} = \sum_{t=50ps}^{T} \left[\left(rmsf_0 * \log\frac{rmsf_0}{rmsf_1}\right) + \left(rmsf_1 * \log\frac{rmsf_1}{rmsf_0}\right)\right] \Big/ T \tag{12}$$

where *rmsf* represents the average root mean square deviation of a given atom over time t. More specifically, the *rmsf* is a directionless root mean square fluctuation sampled over an ensemble of MD runs with similar time slice intervals. Because mutational events in evolution most often replace entire residues, this calculation is more useful if applied to resolution of single amino acids rather than single atoms. Because only the 4 protein backbone atoms (N, Cα, C and O) are homologous between residues, the R groups or side chains are ignored in the calculation and the following equation can be applied. Because the sidechain atoms always attach to this backbone, *rmsf* still indirectly samples the dynamic effect of amino acid sidechain replacement as they are still present in the simulation.

$$RE_{functional} = \sum_{t=50ps}^{T} \sum_{i=N,C,C_\alpha,O}^{4} \left[\left(rmsf_0 * \log\frac{rmsf_0}{rmsf_1}\right) + \left(rmsf_1 * \log\frac{rmsf_1}{rmsf_0}\right)\right] /8 \Big/ T \tag{13}$$

When applied to the entire protein and standardized for protein size, relative entropy can be calculated as

$$RE_{functional} = \sum_{i=1}^{R} \sum_{t=50ps}^{T} \sum_{j=N,C,C_\alpha,O}^{4} \left[\left(rmsf_0 * \log\frac{rmsf_0}{rmsf_1}\right) + \left(rmsf_1 * \log\frac{rmsf_1}{rmsf_0}\right)\right] / 8 \Big/ T/R \tag{14}$$

where R is the number of amino acid residues in the protein.

Note that the relative entropy calculation is made symmetric by averaging with the result obtained in the opposite direction (i.e. swapping state 0 and state 1). It can also be signed in accordance with whether *rmsf* values increase or decrease between states 0 and 1. Also note that any dynamic feature of atom motion can be used in place of *rmsf*. For example, protein shape changes over time could be captured with vector trajectory analysis and analyzed in conjunction with *rmsf*.

B) Relative entropy between conserved dynamic states of variants

When experimenting with *in silico* site-directed mutagenesis or with differently docked small molecule ligands, it may be more useful to compare the impacts of these variants on the functionally conserved dynamics (defined previously) rather than features like *rmsf* representing on the total dynamics. This is important as not all dynamics may be related to the protein function. The relative entropy between the local self-correlation in learning performance (i.e. functional dynamic information contributing to learning performance of the two validation runs) and each variants correlation to the validation run can prove useful for comparing the functional dynamic impacts of each molecular variant.

$$RE_{variant} = \sum_{i=1}^{r} CC_{self} * \log\frac{CC_{self}}{CC_{variant}} \tag{15}$$

Where CC is a canonical correlation analysis of learning performance profiles along length of the protein comparing two identical MD validation runs (i.e. CC_{self}) or a variant to one of the validations (i.e. $CC_{variant}$). Empirical p-values for determining if genetic variants or drug classes differ significantly can be easily generated by randomization tests conducted with bootstrapped frequencies of classification (i.e. learning performance) generated by the time slices of MD variant runs.

C) Mutual information and coordinated dynamics

Mutual information (or trans-information) (5,29) is defined as the relative entropy between the joint distribution or co-occurrence of two events A and B and the product distributions of their individual occurrence. Thus

$$MI = P(A,B) * \log \frac{P(A,B)}{P(A) * P(B)} \tag{16}$$

where P is the probability of two events A and B that may or may not occur independently

If we were to consider the co-occurrence of the functional dynamic classification states over time on any two sites of the protein, we can use mutual information to identify concerted dynamic functions between functional protein sites. This can potentially be very useful for the study of dynamics of proteins with bilateral binding symmetry (30,31) or allosteric regulation (32). If the dynamics of two neighboring sites or even two distant sites are locked together in time, they will exhibit higher degrees of mutual information. Therefore, following the machine learning-based functional dynamic analysis outlined above, a mutual information (MI) matrix calculated over all functionally conserved sites can identify regions where neighboring and/or more distant sites exert influence over each other's dynamic function.

The mutual information between two functional dynamic classification states (0=unbound and 1=bound) at any two given residue sites (r1 and r2) calculated over t time intervals can be calculated as

$$MI = \frac{1}{2}\left[\sum_{t=50ps}^{T} f(0_{r1}, 0_{r2}) * \log \frac{f(0_{r1}, 0_{r2})}{f(0_{r1}) * f(0_{r2})} + \sum_{t=50ps}^{T} f(1_{r1}, 1_{r2}) * \log \frac{f(1_{r1}, 1_{r2})}{f(1_{r1}) * f(1_{r2})}\right] \tag{17}$$

A case study of the functionally conserved dynamic analysis (FCDA) of the RNA polymerase II Rbp4 subunit

Here we demonstrate the efficacy of our proposed method of functionally conserved dynamic analysis (FCDA), using molecular dynamic simulations of the bound and unbound states of Rbd4, the fourth largest subunit of RNA polymerase II. Rbd4 functionally binds Rdb7 within the RNA pol II structure (Figure 2A) and has well characterized conserved regions in the N and C terminus connected by a non-conserved linker region (Figure 2B). Functional studies in yeast mutants have established clear links between many functional yeast phenotypes and deletions in these regions. We used our FCDA protocol above to examine whether functionally conserved dynamics would be associated with the functionally conserved sequence regions previously established by the work of (21). For comparison, we compared

this protein-protein interaction to a DNA-protein interaction in the form of human TATA binding protein (TBP) bound to DNA (33). The software implementation of the FCDA protocol is recently published by our group (15) and software pipeline was named 'maxDemon' in reference to Maxwell's mythological demon that observes the motion of atoms to achieve information gain.

PDB structure preparation and MD simulation protocols

Structures of the human and yeast Rbp4 bound to Rbp7 were obtained from the Protein Data Bank (PDB). These were PDB ID: 2c35 and 1y14 respectively (18–20). We similarly obtained human TBP structure (PDB: 1cdw). Large ensembles of graphic processing unit (GPU) accelerated molecular dynamic simulations were prepared and conducted using the particle mesh Ewald method employed by pmemd.cuda in Amber18 (11,34) via the DROIDS v3.8 interface (Detecting Relative Outlier Impacts in Dynamic Simulation) (15,16). Simulations were run on a Linux Mint 19 operating system mounting two Nvidia Titan Xp graphics processors. Explicitly solvated protein systems were prepared using tLeAP (Ambertools18) using the ff14SB protein force field (35). For TBP, we also loaded the DNA.OL15 force field (36). Solvation was generated using the Tip3P water model (37) in a 12nm octahedral water box and subsequent charge neutralization with Na+ and Cl- ions. After energy minimization, heating to 300K, and 10ns equilibration, an ensemble of 200 MD production runs each lasting 0.5 ns of time were created for both bound and unbound Rbp4 in both the yeast and human MD analysis. Each MD production run was preceded by a single random length short spacing run selected from a range of 0 to 0.25ns to mitigate the effect of chaos (i.e. sensitively to initial conditions) in the driving ensemble differences in the MD production runs (15). All MD was conducted using an Andersen thermostat (38) under constant pressure of one atmosphere. Root mean square atom fluctuations (*rmsf*) were calculated using the atomicfluct function in CPPTRAJ (24). All color maps on protein structures were produced in UCSF Chimera (39).

Functionally Conserved Dynamic Analysis (FCDA) of bound and unbound states protein subunits within RNA Polymerase II and DNA-bound and unbound states of TATA binding protein

The signed symmetric relative entropy between the distributions of atom fluctuation (i.e. root mean square fluctuation or *rmsf* taken from 0.01 ns time slices of total MD simulation time) on bound and unbound protein states were computed using our program DROIDS and color mapped to the protein backbone with individual amino acid resolution to the bound structures using a temperature scale (i.e. where red is hot or amplified fluctuation and blue is cold or dampened fluctuation). The reference state of the protein is unbound while the query state is bound. Therefore, this pairwise comparison is used to represent the functional impact of Rbp7 binding on the Rbp4 protein's normal unbound motion, where it is expected that contacts would typically dampen the fluctuation of atoms around the binding interface to measurable degree. In the case of TATA binding protein (TBP), the functional dynamic comparison was comprised of DNA bound to TBP and unbound TBP, again where it is expected that DNA binding should dampen the atom fluctuation to some degree. Functionally conserved dynamics were determined using the method outlined previously with a stacked machine learning model that included 7 classification methods to classify *rmsf* time series values into either 'bound' or 'unbound' classifications. The classifiers included K nearest neighbors, naïve Bayes, linear discriminant function, quadratic discriminant function, support vector machine, random forest and adaptive boosting. Significant canonical correlations in the 7 local machine learner performance profiles were determined using Wilk's lambda and color mapped to dark gray regions on the bound Rbp4 and TBP structures. A

mutual information matrix to map coordinated classifications of dynamic states over time between all amino acid residue pairs was also calculated. All software used to produce all statistical and graphical output is part of a comprehensive package DROIDS/maxDemon version 3.8 available at GitHub (https://github.com/gbabbitt/DROIDS-3.0-comparative-protein-dynamics) and our software website (https://people.rit.edu/gabsbi/) and its domain (http://proteindynamics.net).

Results

FCDA case study in RNA Polymerase II

The relative entropy in residue backbone *rmsf* between Rbp7 bound Rbp4 and unbound Rbp4 dynamic states (i.e. eqn. 13) revealed common patterns of dampening in atom fluctuation (i.e. rapid harmonic motion) in both human and yeast. In accordance with the prediction from the functional work of Sampath et al (2003), both of our *in silico* comparative dynamics experiments show more functional dampening of atom motion in the conserved N and C terminal regions with less dampened motion in the non-conserved central linker regions (Figure 3). Note that the linker region is much larger in yeast, and our patterns of dampened atom fluctuation reflect this as well (Figure 3B). The most dampened motion coincides consistently with Rbp4's contact surface with Rbp7 (Figure 4) which proportionally comprises a much larger fraction of the protein in human Rbp4 (Figure 4A-C) than yeast (Figure 4 D-F) and best compared in the side profiles (Figure 4B and 4E). The machine learning profiles (i.e. frequency term in eqn. 8) on the Rbp7 bound Rbp4 MD validation runs (Figure 5A/5C) indicate that motions indicative of functional binding were very well classified (i.e. recognized by the learning model) in the same regions showing more dampened motion or large negative relative entropy in Figures 3. This indicates that our machine learning classifier was able to recognized differences in *rmsf* related to functional binding between the two subunits of RNA Pol II. Significant Wilk's lambda for canonical correlations indicate regions of functionally conserved dynamics that correspond very well with regions of dampened *rmsf* upon binding (Figure 3) and the functionally conserved regions associated with the Rbp4/7 interface identified by (21)(Figure 5B and 5 D). The yeast linker region which is largely non-conserved at the sequence level according to Sampath (21) shows very little conserved dynamics according to our method (Figure 5D). The human Rbp4, which lacks the large linker region, exhibits much more conserved dynamics throughout the protein (Figure 5B). The mapping of functionally significantly conserved dynamic regions to the Rbp4 structures also indicates that Rbp4 linker regions more distant from the Rbp7 interface are far less likely to be classified by learners as having a role in functional binding than Rbp4 region interacting more directly with the binding interface (Figure 6). Significance tests on the relative entropy of *rmsf* and R values derived from canonical correlations of machine learning profiles correspond roughly with the 10 period moving average of sequence-based Shannon information (r = 0.487/0.484 for human/yeast Rbp4) indicating that much of the Rbp4 protein sequence information maintained across distant taxonomic groups is contributing to the maintenance of functionally conserved dynamic motion related to Rbp7 binding (Supplemental Figure 1). Lastly, mutual information matrices (eqn. 17, Figure 7) indicate that time coordinated dynamic motions between amino acid residues are quite prevalent in the non-conserved linker regions of Rbp4, and that this is more readily apparent in the central linker region of yeast Rbp4 (Figure 7C) than in human Rbp4 (Figure 7B) or random control (Figure 7A).

FCDA case study in TATA binding protein

We conducted an identical FCDA analysis comparing the dynamics of DNA bound TBP to its unbound form (Figure 8). Relative entropy calculations indicated where the major groove contacts of TBP are (Figure 8A and 8B). The machine learning analysis of multiple validation runs on DNA bound TBP indicated very high and significant canonical correlations in machine learning recognition of binding state across almost the entire structure of TBP (Figure 8C). A mutual information matrix also indicated very high levels of time coincidental classifications of indicating very strong coordination of functional dynamic states across the whole structure of TBP (Figure 8D). As the DNA binding partner of TBP is much more rigid than the Rbp7 protein binding partner to Rbp4, there is consequently much more mutual information between the functional classifications of dynamics between different protein sites.

Discussion

We have introduced an information theoretical formalism and a general protocol for adapting machine learning classification applied to molecular dynamics simulation towards the task of identifying functionally conserved and coordinated motions in proteins. While molecular dynamics simulations are traditionally computationally heavy, this method can now easily be accelerated on modern graphics processors and potentially allows the molecular biologist to step beyond the functional abstraction of information contained in sequences, and to directly analyze information content in the functional motion of proteins in accurate biophysical simulations. This allows for the definition of functional conservation within the known context of specific function of the proteins under investigation. Because machine learning classification binary digitizes molecular dynamic output, it allows information theory to be analytically applied in a variety of useful ways for discerning random thermal events (i.e. entropy) from non-random motion (i.e. function) in the context of protein interactions with other proteins, nucleic acids or small molecules. This protocol can now be implemented in the latest releases of our labs software website https://people.rit.edu/gabsbi/ or http://proteindynamics.net and published software notes (15,16). The information theoretics incorporated in the latest release of DROIDS+maxDemon version 3.8 include all the methods introduced here, including the comparison of *rmsf* between functional states via relative entropy, the identification of functionally conserved dynamics (FCDA) in new MD runs, the comparison of relative impacts of variants on conserved dynamics via relative entropy, and the identification of coordinated dynamic motions across sites on the protein via the calculation of a mutual information matrix.

We have demonstrated the efficacy of this information theory based machine learning approach to identifying conserved dynamics by using simulations of functional binding in the well-understood Rbp4/7 subunit interaction in RNA polymerase II. Our method can quantify differences in local dynamics using relative entropy calculations and can recognize and map regions of local conserved dynamics by identifying non-random local learning performance peaks indicative of functional training states that defined the relative entropy calculations. Most importantly, our results computationally confirm the main functionally conserved relationships of the N and C terminal regions investigated through prior sequence analysis and site-directed mutagenesis studies (21–23,40). We also demonstrate that the pairwise calculation of mutual information can be used to identify time coordinated dynamics even in regions where sequences and binding dynamics are not conserved (e.g. linker regions). Site-directed mutagenesis studies on Rbp4 have indicated that some mutations in non-conserved regions can have functional effects on phenotype. Linker regions that are not conserved at the sequence level can still

have important physical function imparted by structural rigidity. By isolating time-coordinated molecular dynamics between sites, our mutual information matrices capture this feature of linker regions very well.

Our analyses of TATA binding protein (TBP) also showed the potential effectiveness of our computational approach in illuminating important functional dynamics. First of all, the relative entropy calculation on atom fluctuation *(rmsf)* highlights the locations of major groove contacts very well, and quantifies the binding effect as much larger than the rest of the protein. Second, this analysis also indicates that the whole structure of TBP is dampened in its motion during binding even outside of these major groove contacts, and this dampening is largely a function of any given protein atom's proximity to the DNA. The realization that nearly all of the TBP protein is functionally involved in binding is also confirmed by the machine learning performance analysis and identifies almost the whole TBP structure as exhibiting significantly functionally conserved protein dynamics. Lastly, the mutual information matrix of all pairwise amino acid residues indicates very high levels of time coincident machine classification of functional dynamics states in the validation runs, even across very distant parts of the protein. This suggests that, unlike Rbp4, which binds the softer matter of Rbp7 and in two distinct regions separated by a linker section, the much higher rigidity of DNA is likely involved in coordinating long range dynamic correlations in concerted motion across the whole TBP structure when it is in its DNA bound state. Thus, the lack of a designated non-conserved linker region in TBP, combined with the rigidity of its DNA binding partner, creates very large degree of mutual dynamics when compared to Rbp4, which only seems to show mutual dynamics in its linker region. Our method captures this aspect of shared dynamics very well.

Molecular biologists typically investigate molecular function by discovering DNA/protein sequences and then conducting a series of expensive laboratory studies that probe function via experimental manipulation of these sequence changes. While protein structures are often characterized as part of this pipeline, very rarely is computer simulation of dynamic motion ever employed. Methods of biophysical simulation of protein dynamics have been under continuous development for many decades. However, most of the questions addressed through simulation are physical rather than comparative in nature. The lack of a common comparative method for molecular dynamics data has been problematic. Many reviews of best practices in molecular dynamics often warn users of the danger of comparisons based upon inadequate sampling, but seldom offer proper statistical framework to mitigate this problem (41,42). Recently, our lab has developed a platform for statistical comparison and machine learning applications to molecular dynamics (15,16). Here, we have broadened this application under an information theoretical framework to comparative dynamics that has the potential to illuminate the functional consequences of molecular changes at an all-atom resolution that cannot be achieved with traditional molecular biology. Our computational protocol helps to solve an important problem of scale in molecular biology; the investigation of functional molecular motions at the very rapid time scales and very small size scales at which it occurs. It also offers the ability to connect these very rapid short time scales to the very long time scales of molecular evolution. Frameworks where machine learning is applied to molecular dynamic simulation can allow proper investigation of this complex and largely invisible process involving time scales that are not easy for humans to directly observe.

An understanding of molecular biology that is mainly based upon sequence and structure, while ignoring dynamics, will create data sets with large underlying latent variability and will always

potentially hinder interpretations of lab-based experiments. The growing complexity of many sequence-based assay technologies in molecular biology are allowing many new questions to be asked but are also creating large complex datasets with a potential for underlying latent and systematic bias. Computational experiments conducted in simulation and properly analyzed from the standpoint of thermodynamic information and entropy provide insight into the dynamic nature of protein interaction at all-atom resolution can be quite valuable in supplementing the knowledge derived from current high-throughput sequence-based technologies.

Figure Captions

Figure 1. A visual analogy of Functionally Conserved Dynamic Analysis (FCDA) applied to a hypothetical ensemble of protein simulations in two functional states (RNA POL II subunit Rbp4 bound to Rbp7 and Rbp4 in its unbound state). (A) Functional motion is evolutionarily conserved in non-random state and can be recognized by machine learning classification. (B) Local learning performance increases in functional regions of the protein and (C) can be mathematically retrieved by canonical correlation analysis of independent molecular dynamic validation runs.

Figure 2. The functional binding of Rbp4 within RNA Pol II. (A) The 12 unit RNA Pol II (PDB: 1wcm) highlighting the relative positions of subunits Rbp4 and Rbp7 (PDB: 1y14) where the comparative molecular dynamics analysis of Rbp4 implemented. (B) The 10 period smoothed Shannon information content for the multiple sequence alignment of Rbp4 from Sampath et al. 2003. The alignment includes Saccharomyces cerevisiae, Schizosaccharomyces pombe, Arabidopsis thaliana, Drosophila melanogaster, and Homo sapiens.

Figure 3. Quantifying the functional binding interaction of Rbp4/7 within RNA Pol II with a signed symmetric Kullback-Leibler distance (i.e. relative entropy of eqn. 13) comparing root mean square fluctuation (rmsf) of protein backbone for Rbp7 bound Rbp4 and unbound Rbp4 subunits of RNA Pol II in both (A) *Homo sapiens* and (B) yeast *Saccharomyces cerevisiae*. Note that more negative KL divergences or dFLUX indicates a stronger functional binding interaction between Rbp7 and Rbp4.

Figure 4. Imaging the functional binding interaction of Rbp4/7 within RNA Pol II with color mapping of signed symmetric Kullback-Leibler distance (i.e. relative entropy of eqn. 13) from Figure 3 in (A-C) front, side and back profiles of human Rbp4 and (D-F) the identical profiles for yeast, *S cerevisiae*. Note: Rbp4 is the dFLUX surface mapped structure while Rbp7 is indicated with lavender ribbon.

Figure 5. Functionally conserved dynamics of Rbp4 identified via local canonical correlation analysis of multi-machine learning performance profiles in (A-B) humans and (C-D) yeast, *S. cerevisiae*. Machine learning performance profiles (frequency term in eqn. 8) were generated in two independent molecular dynamics validation runs on Rbp7 bound Rbp4 (red and blue lines) using multiple classification methods including K-nearest neighbor, naïve Bayes, linear discriminant analysis, quadratic discriminant analysis, support vector machine, random forest and adaptive boosting (A and C). The local R value for the canonical correlation is shown over an 18 residue sliding window (orange line in B and D) and its regions with significant correlation (i.e. determined via Wilk's lambda) is demarcated using gray shading under the green line.

Figure 6. **Mapping of functionally conserved dynamics of Rbp4 identified via local canonical correlation analysis of multi-machine learning performance profiles in (A-C) humans and (D-F) yeast, *S. cerevisiae*.** Local regions of Rbp4 with significant functionally conserved dynamic interaction (eqn. 9) with Rbp7 (blue ribbon) are mapped in dark gray.

Figure 7. **Mutual information matrices for (A) a null example with very little dynamic differences to learn and examples of Rbp7 bound Rbp4 in (B) humans and (C) yeast, *S. cerevisiae*.** Black tiles indicate no matching of classification of functional dynamic states between two given amino acid residues (i.e. no mutual information from equation 17) while white indicates complete matching of dynamic states (i.e. total mutual information from equation 17). The machine learning classifier in these examples was a parameter tuned support vector machine with polynomial kernel.

Figure 8. **Functionally conserved dynamic analysis (FCDA) and the associated information theoretics for human TATA binding protein in DNA-bound and unbound functional states (PDB: 1cdw)** showing (A) negative relative entropy associated with binding that is also (B) color mapped to the DNA bound protein structure. Note that blue regions indicate dampened protein atom fluctuations associated with functional binding. Yellow arrows indicate loop structures that interact with DNA major groove. (C) Functionally conserved dynamics regions are determined via significant Wilk's lambda value ($p<0.01$) determined from canonical correlation analysis of machine learning validation profiles within 30 residue sliding window. Conserved dynamics are indicated by dark grey regions in both the plot and the structure (inset image). (D) A mutual information matrix indicating where functional dynamic classifications coincide over time is also shown. Lighter tiles indicate high degree of time dependent mutual information in functional dynamic classifications between amino acid pairs.

Supplemental Figure A. Co-occurrence of sequence-based Shannon information content and functionally conserved molecular dynamics of Rbp4/7 interaction within RNA Pol II. The Shannon information (black lines) was calculated using the multiple sequence alignment of Sampath et al. 2003. The metrics of functionally conserved dynamics (R values) and maximum statistical differences in the empirical cumulative distributions of root mean square fluctuation of Rbp4 in bound versus unbound states (D values) are shown as colored lines for comparison. Note that zero values indicate large gaps in the main alignment where yeast linker sequence is not represented in human sequence.

Acknowledgements

I thank and acknowledge Dr. Ernest P. Fokoue, Dr. Miranda Lynch and Dr. George M Thurston for very helpful conversations regarding this subject.

Funding

There is no funding source to report.

References

1. Davies PCW. The Demon in the Machine [Internet]. [cited 2020 May 11]. Available from: https://press.uchicago.edu/ucp/books/book/chicago/D/bo45084244.html

2. Davies PCW, Rieper E, Tuszynski JA. Self-organization and entropy reduction in a living cell. Biosystems. 2013 Jan;111(1):1–10.

3. Walker SI, Davies PCW. The algorithmic origins of life. J R Soc Interface. 2013 Feb;10(79):20120869.

4. Schrödinger E, Penrose R. What is Life?: With Mind and Matter and Autobiographical Sketches [Internet]. Cambridge Core. Cambridge University Press; 1992 [cited 2020 May 11]. Available from: /core/books/what-is-life/A876185F2DB06FF5C2CC67C9A60DAD7F

5. Shannon CE. A mathematical theory of communication. Bell Syst Tech J. 1948 Jul;27(3):379–423.

6. Cobb M. 1953: when genes became “information.” Cell. 2013 Apr 25;153(3):503–6.

7. Schneider TD, Stormo GD, Gold L, Ehrenfeucht A. Information content of binding sites on nucleotide sequences. J Mol Biol. 1986 Apr 5;188(3):415–31.

8. Schneider TD, Stephens RM. Sequence logos: a new way to display consensus sequences. Nucleic Acids Res. 1990 Oct 25;18(20):6097–100.

9. Babbitt GA, Coppola EE, Alawad MA, Hudson AO. Can all heritable biology really be reduced to a single dimension? Gene. 2016 Mar 10;578(2):162–8.

10. Götz AW, Williamson MJ, Xu D, Poole D, Le Grand S, Walker RC. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born. J Chem Theory Comput. 2012 May 8;8(5):1542–55.

11. Salomon-Ferrer R, Götz AW, Poole D, Le Grand S, Walker RC. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 2. Explicit Solvent Particle Mesh Ewald. J Chem Theory Comput. 2013 Sep 10;9(9):3878–88.

12. Babbitt GA, Lynch ML, McCoy M, Fokoue EP, Hudson AO. Function and evolution of B-Raf loop dynamics relevant to cancer recurrence under drug inhibition. bioRxiv. 2020 Jan 13;2020.01.13.904052.

13. Babbitt GA, Coppola EE, Mortensen JS, Ekeren PX, Viola C, Goldblatt D, et al. Triplet-Based Codon Organization Optimizes the Impact of Synonymous Mutation on Nucleic Acid Molecular Dynamics. J Mol Evol. 2018;86(2):91–102.

14. Babbitt GA, Alawad MA, Schulze KV, Hudson AO. Synonymous codon bias and functional constraint on GC3-related DNA backbone dynamics in the prokaryotic nucleoid. Nucleic Acids Res. 2014;42(17):10915–26.

15. Babbitt GA, Fokoue EP, Evans JR, Diller KI, Adams LE. DROIDS 3.0-Detecting Genetic and Drug Class Variant Impact on Conserved Protein Binding Dynamics. Biophys J. 2020 Feb 4;118(3):541–51.

16. Babbitt GA, Mortensen JS, Coppola EE, Adams LE, Liao JK. DROIDS 1.20: A GUI-Based Pipeline for GPU-Accelerated Comparative Protein Dynamics. Biophys J. 2018 13;114(5):1009–17.

17. Bakan A, Dutta A, Mao W, Liu Y, Chennubhotla C, Lezon TR, et al. Evol and ProDy for bridging protein sequence evolution and structural dynamics. Bioinformatics. 2014 Sep 15;30(18):2681–3.

18. Armache K-J, Mitterweger S, Meinhart A, Cramer P. Structures of complete RNA polymerase II and its subcomplex, Rpb4/7. J Biol Chem. 2005 Feb 25;280(8):7131–4.

19. Armache K-J, Kettenberger H, Cramer P. Architecture of initiation-competent 12-subunit RNA polymerase II. Proc Natl Acad Sci U S A. 2003 Jun 10;100(12):6964–8.

20. Meka H, Werner F, Cordell SC, Onesti S, Brick P. Crystal structure and RNA binding of the Rpb4/Rpb7 subunits of human RNA polymerase II. Nucleic Acids Res. 2005;33(19):6435–44.

21. Sampath V, Rekha N, Srinivasan N, Sadhale P. The Conserved and Non-conserved Regions of Rpb4 Are Involved in Multiple Phenotypes in Saccharomyces cerevisiae. J Biol Chem. 2003 Dec 19;278(51):51566–76.

22. Sharma N, Marguerat S, Mehta S, Watt S, Bähler J. The fission yeast Rpb4 subunit of RNA polymerase II plays a specialized role in cell separation. Mol Genet Genomics. 2006 Dec;276(6):545–54.

23. Zhao Y, Li KKC, Ng KP, Ng CH, Lee KAW. The RNA Pol II sub-complex hsRpb4/7 is required for viability of multiple human cell lines. Protein Cell. 2012 Nov 1;3(11):846–54.

24. Roe DR, Cheatham TE. PTRAJ and CPPTRAJ: Software for Processing and Analysis of Molecular Dynamics Trajectory Data. J Chem Theory Comput. 2013 Jul 9;9(7):3084–95.

25. Härdle W, Simar L, editors. Canonical Correlation Analysis. In: Applied Multivariate Statistical Analysis [Internet]. Berlin, Heidelberg: Springer Berlin Heidelberg; 2007. p. 321–30. Available from: https://doi.org/10.1007/978-3-540-72244-1_14

26. Benjamini Y, Hochberg Y. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J R Stat Soc Ser B Methodol. 1995;57(1):289–300.

27. Doolittle WF, Brunet TDP, Linquist S, Gregory TR. Distinguishing between “Function” and “Effect” in Genome Biology. Genome Biol Evol. 2014 May 9;6(5):1234–7.

28. Graur D, Zheng Y, Price N, Azevedo RBR, Zufall RA, Elhaik E. On the immortality of television sets: “function” in the human genome according to the evolution-free gospel of ENCODE. Genome Biol Evol. 2013;5(3):578–90.

29. Kullback S, Leibler RA. On Information and Sufficiency. Ann Math Stat. 1951 Mar;22(1):79–86.

30. Matsunaga Y, Koike R, Ota M, Tame JRH, Kidera A. Influence of Structural Symmetry on Protein Dynamics. PLOS ONE. 2012 Nov 26;7(11):e50011.

31. Schulze B, Sljoka A, Whiteley W. How does symmetry impact the flexibility of proteins? Philos Transact A Math Phys Eng Sci [Internet]. 2014 Feb 13 [cited 2020 May 3];372(2008). Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3871301/

32. Motlagh HN, Wrabl JO, Li J, Hilser VJ. The ensemble nature of allostery. Nature. 2014 Apr 17;508(7496):331–9.

33. Nikolov DB, Chen H, Halay ED, Hoffman A, Roeder RG, Burley SK. Crystal structure of a human TATA box-binding protein/TATA element complex. Proc Natl Acad Sci U S A. 1996 May 14;93(10):4862–7.

34. Case DA, Cheatham TE, Darden T, Gohlke H, Luo R, Merz KM, et al. The Amber biomolecular simulation programs. J Comput Chem. 2005;26(16):1668–88.

35. Maier JA, Martinez C, Kasavajhala K, Wickstrom L, Hauser KE, Simmerling C. ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. J Chem Theory Comput. 2015 Aug 11;11(8):3696–713.

36. Dans PD, Ivani I, Hospital A, Portella G, González C, Orozco M. How accurate are accurate force-fields for B-DNA? Nucleic Acids Res. 2017 Apr 20;45(7):4217–30.

37. Mao Y, Zhang Y. Thermal conductivity, shear viscosity and specific heat of rigid water models. Chem Phys Lett. 2012 Jul 23;542:37–41.

38. Andersen HC. Molecular dynamics simulations at constant pressure and/or temperature. J Chem Phys. 1980 Feb 15;72(4):2384–93.

39. Pettersen EF, Goddard TD, Huang CC, Couch GS, Greenblatt DM, Meng EC, et al. UCSF Chimera--a visualization system for exploratory research and analysis. J Comput Chem. 2004 Oct;25(13):1605–12.

40. Todone F, Brick P, Werner F, Weinzierl ROJ, Onesti S. Structure of an Archaeal Homolog of the Eukaryotic RNA Polymerase II RPB4/RPB7 Complex. Mol Cell. 2001 Nov 21;8(5):1137–43.

41. Braun E, Gilmer J, Mayes HB, Mobley DL, Monroe JI, Prasad S, et al. Best Practices for Foundations in Molecular Simulations [Article v1.0]. Living J Comput Mol Sci [Internet]. 2019 [cited 2020 May 1];1(1). Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6884151/

42. Grossfield A, Patrone PN, Roe DR, Schultz AJ, Siderius D, Zuckerman DM. Best Practices for Quantification of Uncertainty and Sampling Quality in Molecular Simulations [Article v1.0]. Living J Comput Mol Sci. 2018 Oct 27;1(1):5067.

A Think of protein dynamic function like the motions of dancing

functional protein dynamics can act singularly or interactively

e.g. environmental response

e.g. protein signaling

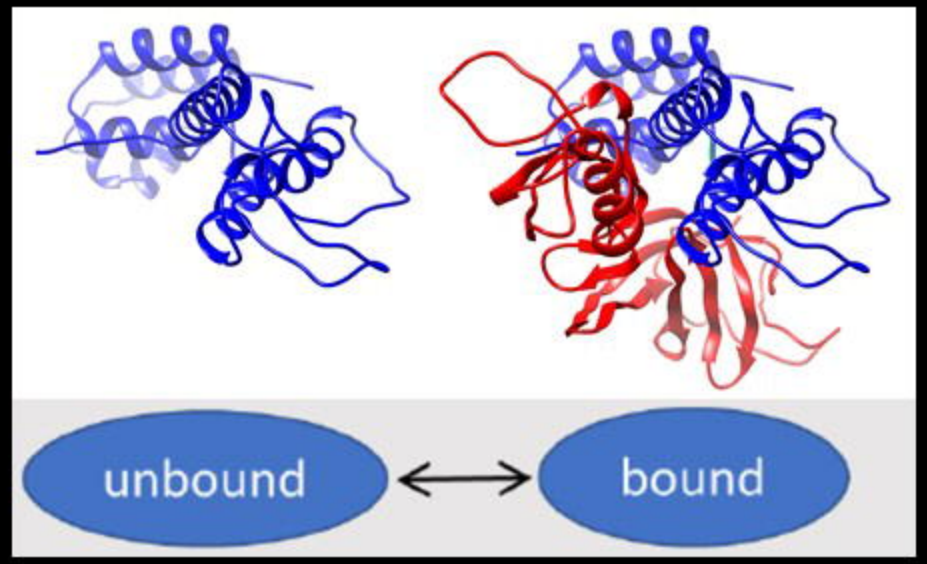

This dynamic motion is

1. Choreographed (i.e. non-random) to a specific task
2. Discernable from random thermal motion
3. Preserved by natural selection over deep time

Functional states can be simulated in molecular dynamics software and resulting data used to train a binary machine learning classifier

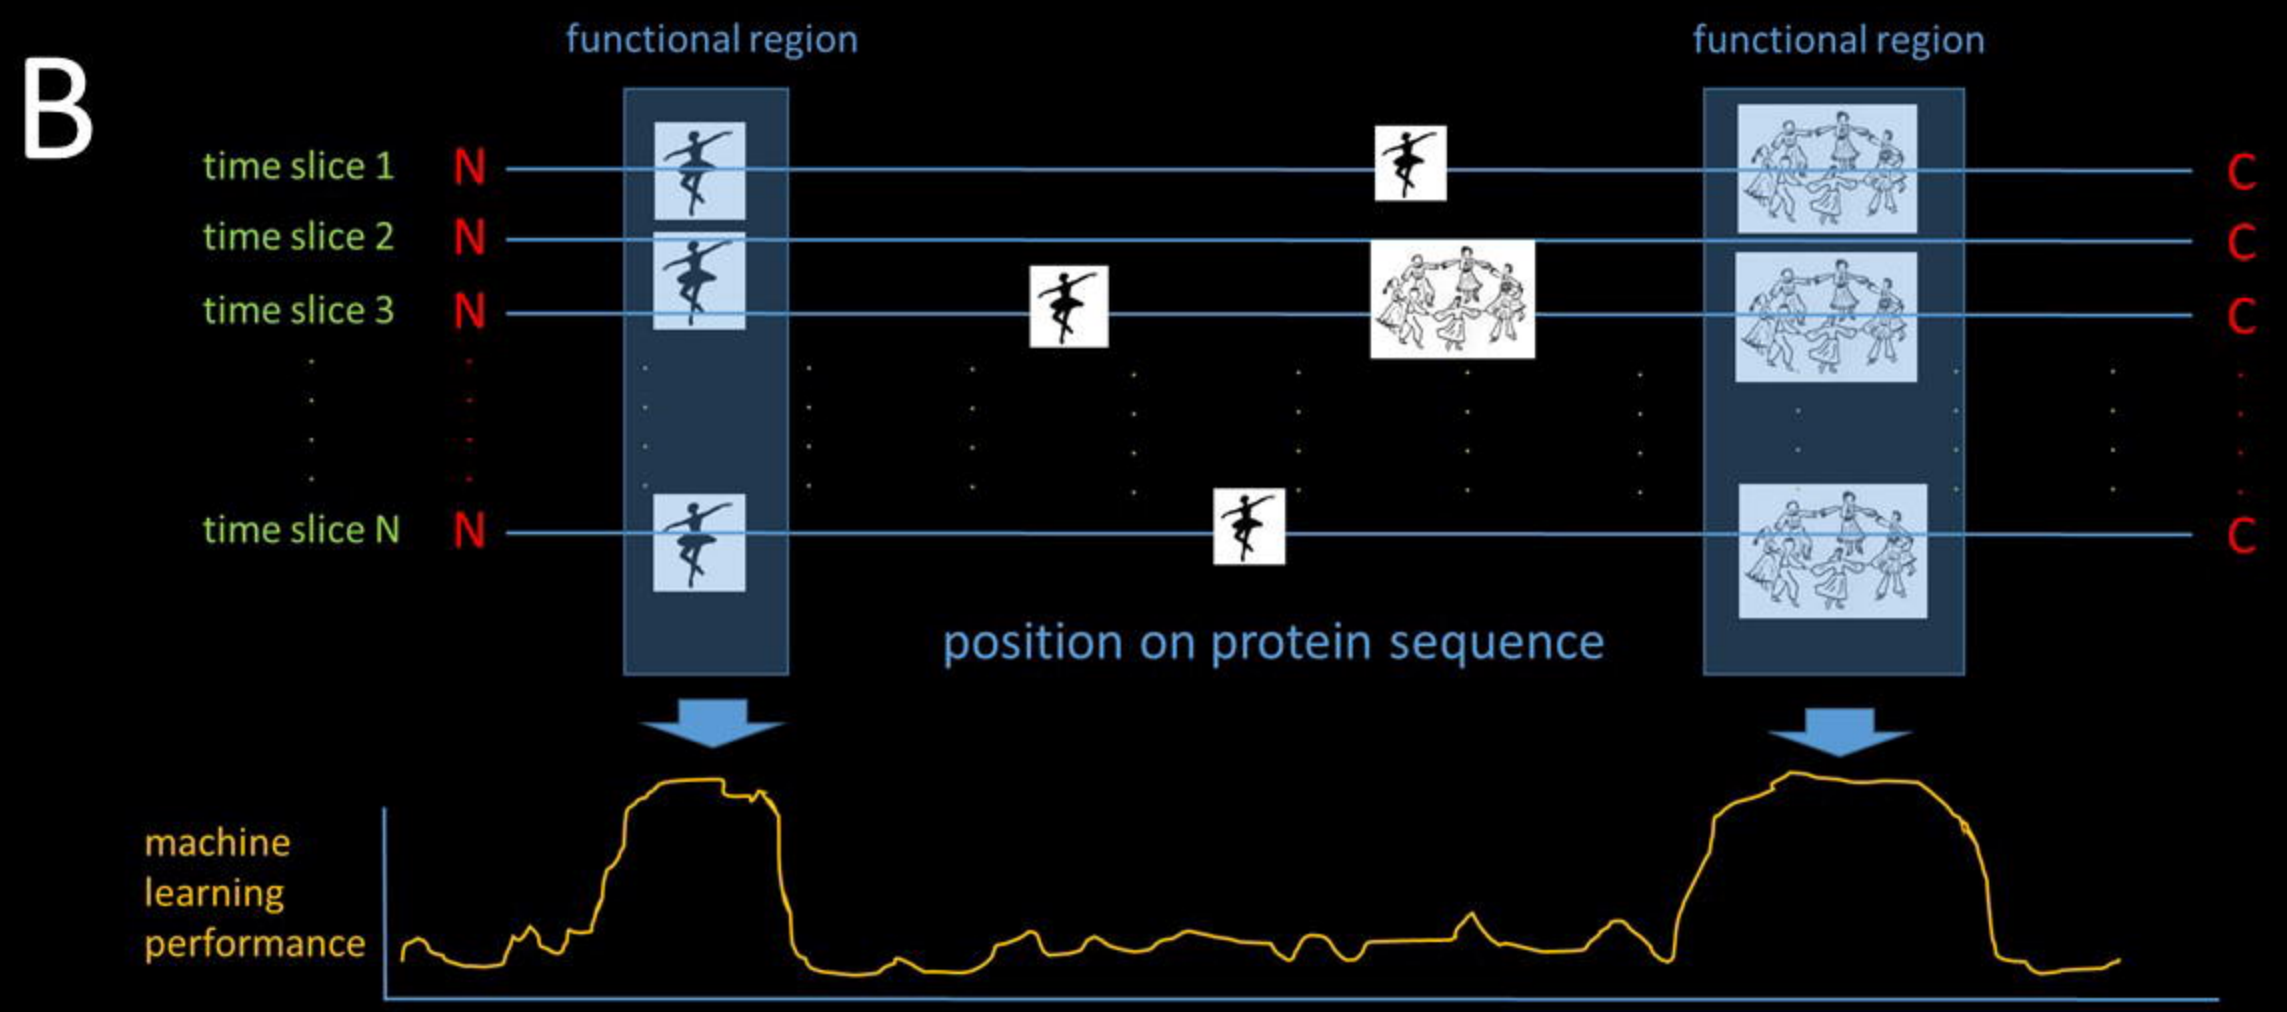

Over many time slices of simulation, conserved dynamic regions will show locally improved trends in machine learning recognition

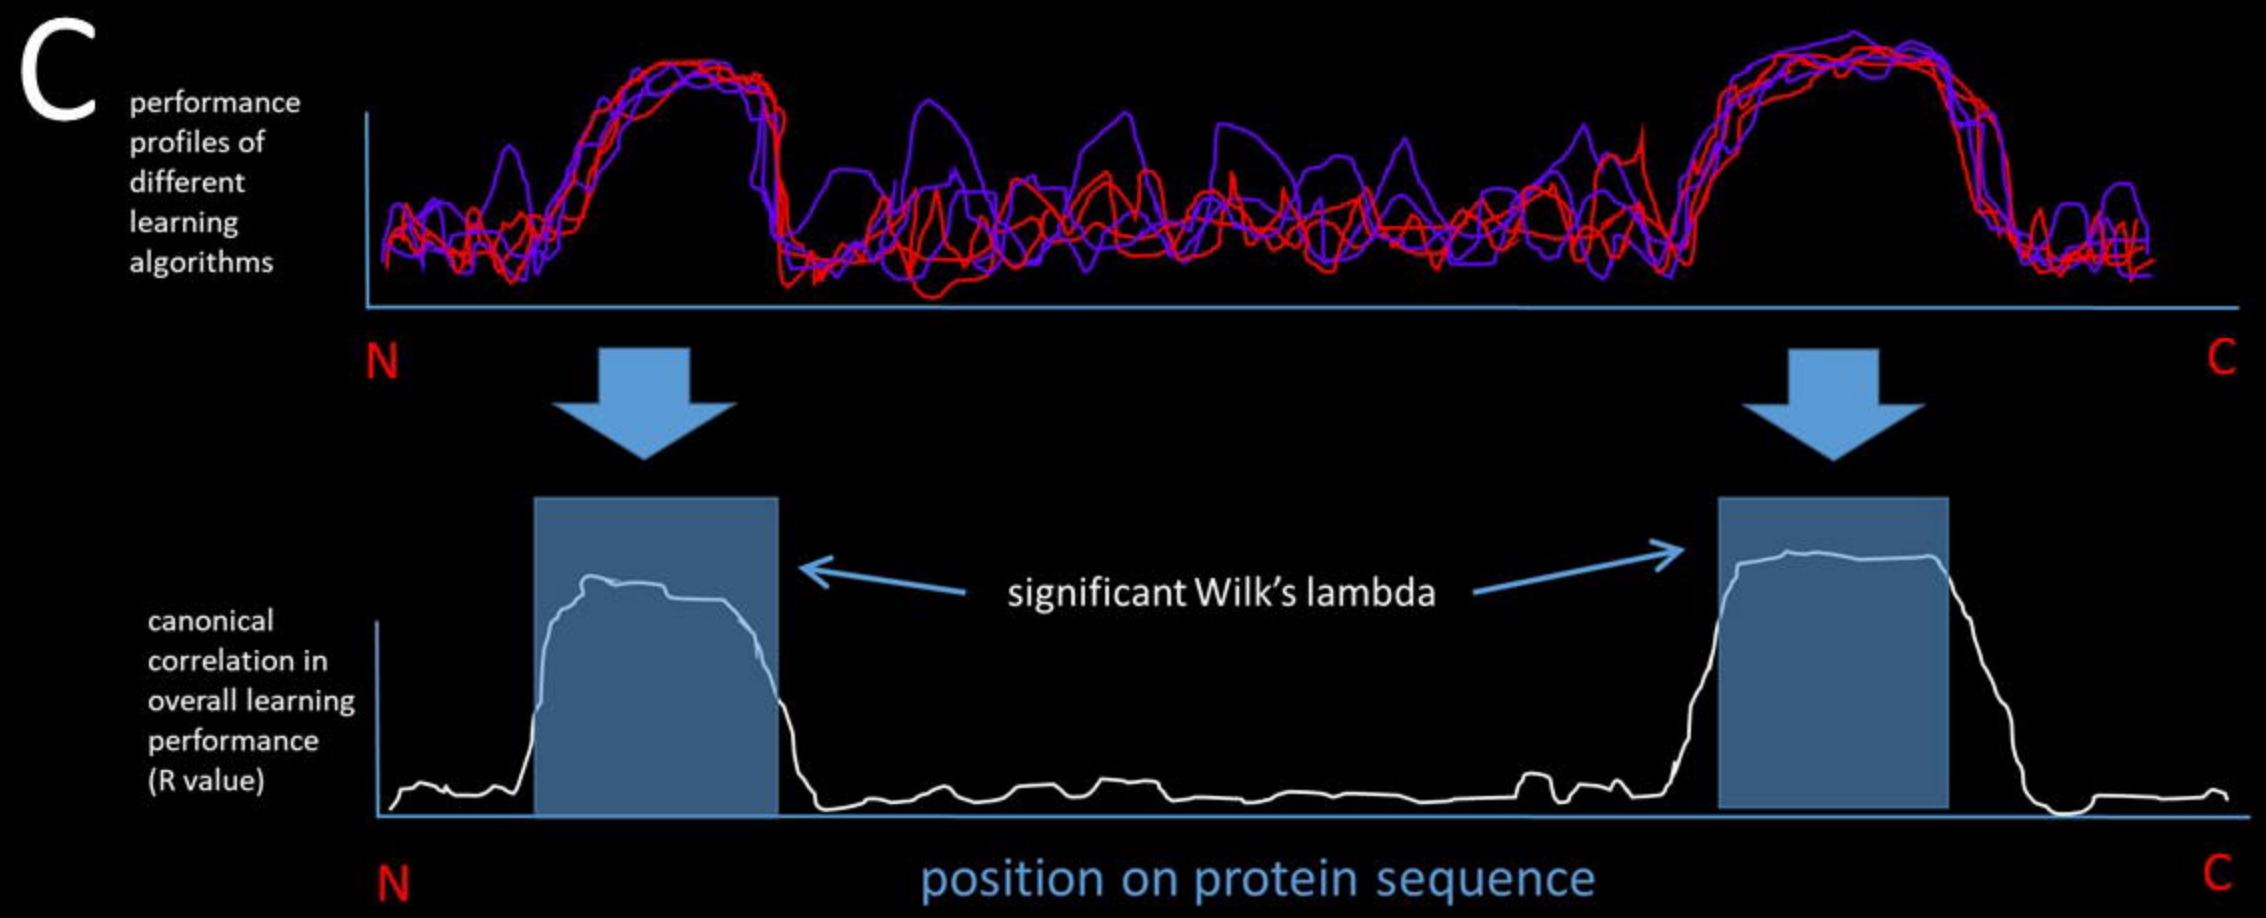

canonical correlation analysis of multi-machine learner performance profiles from two independent validation simulations (red and blue) are used to identify regions with functionally conserved dynamics

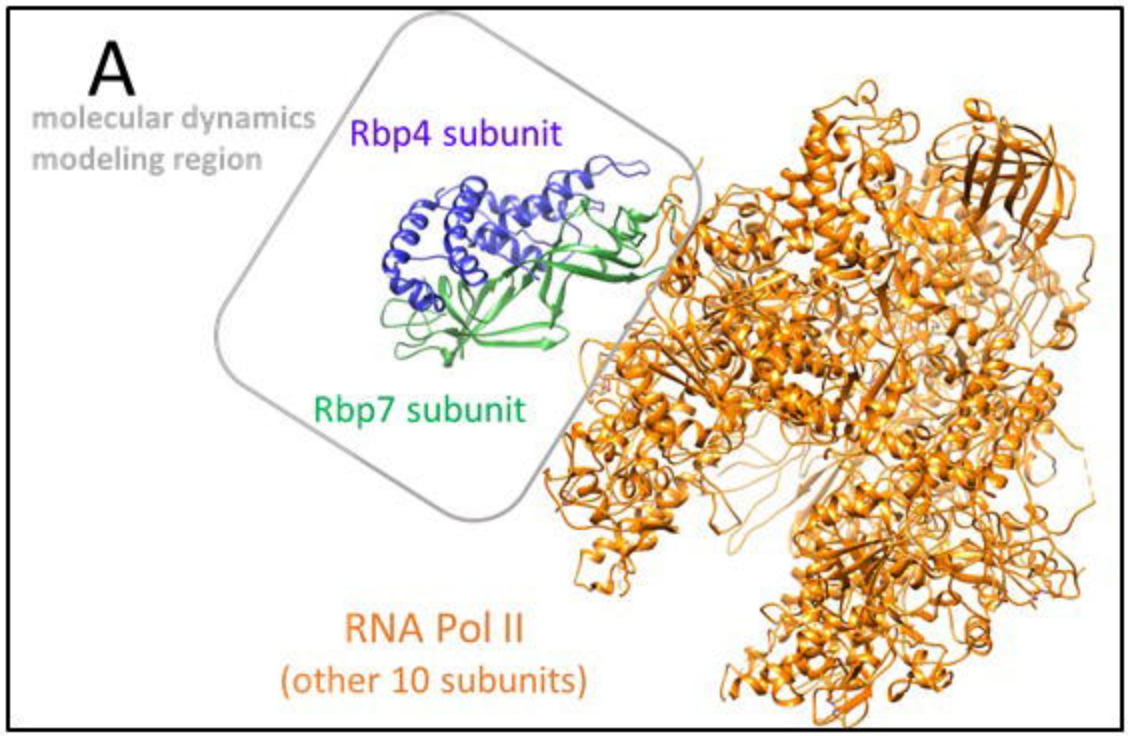
A
molecular dynamics
modeling region
Rbp4 subunit
Rbp7 subunit
RNA Pol II
(other 10 subunits)

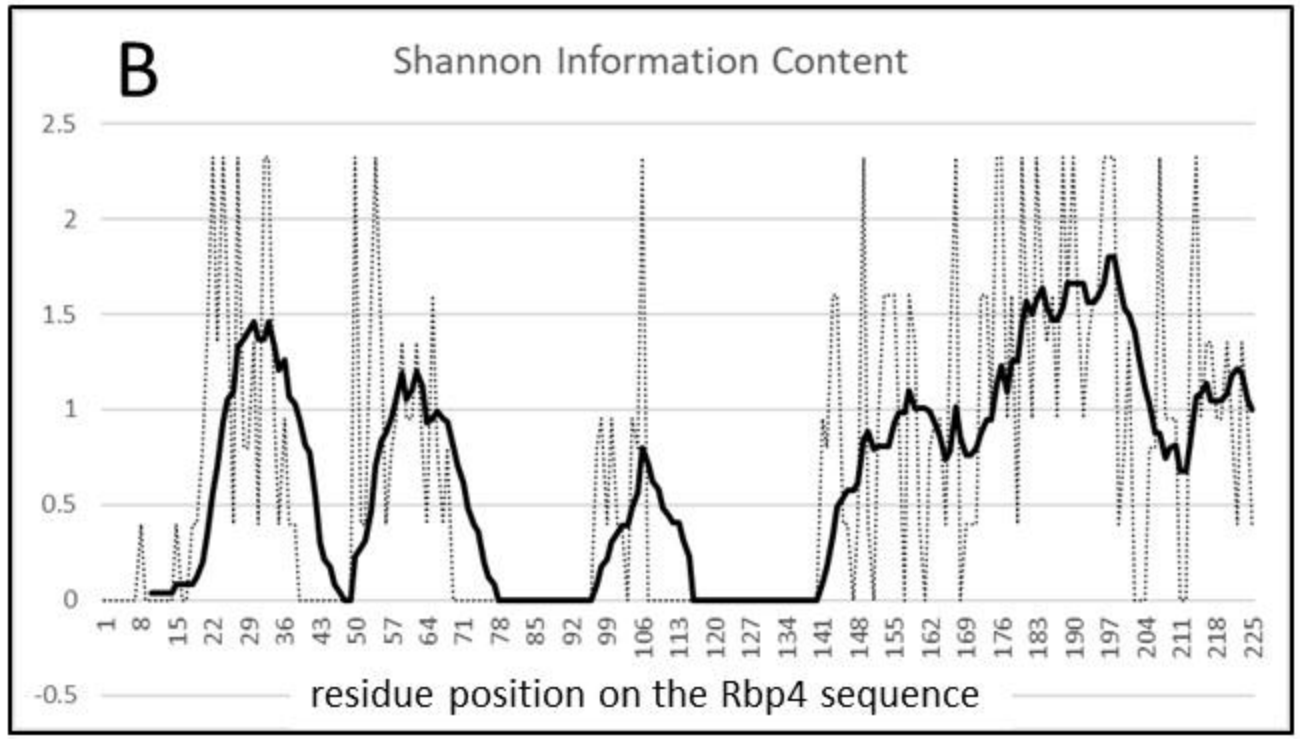
B
Shannon Information Content
2.5
2
1.5
1
0.5
0
-0.5
1 8 15 22 29 36 43 50 57 64 71 78 85 92 99 106 113 120 127 134 141 148 155 162 169 176 183 190 197 204 211 218 225
residue position on the Rbp4 sequence

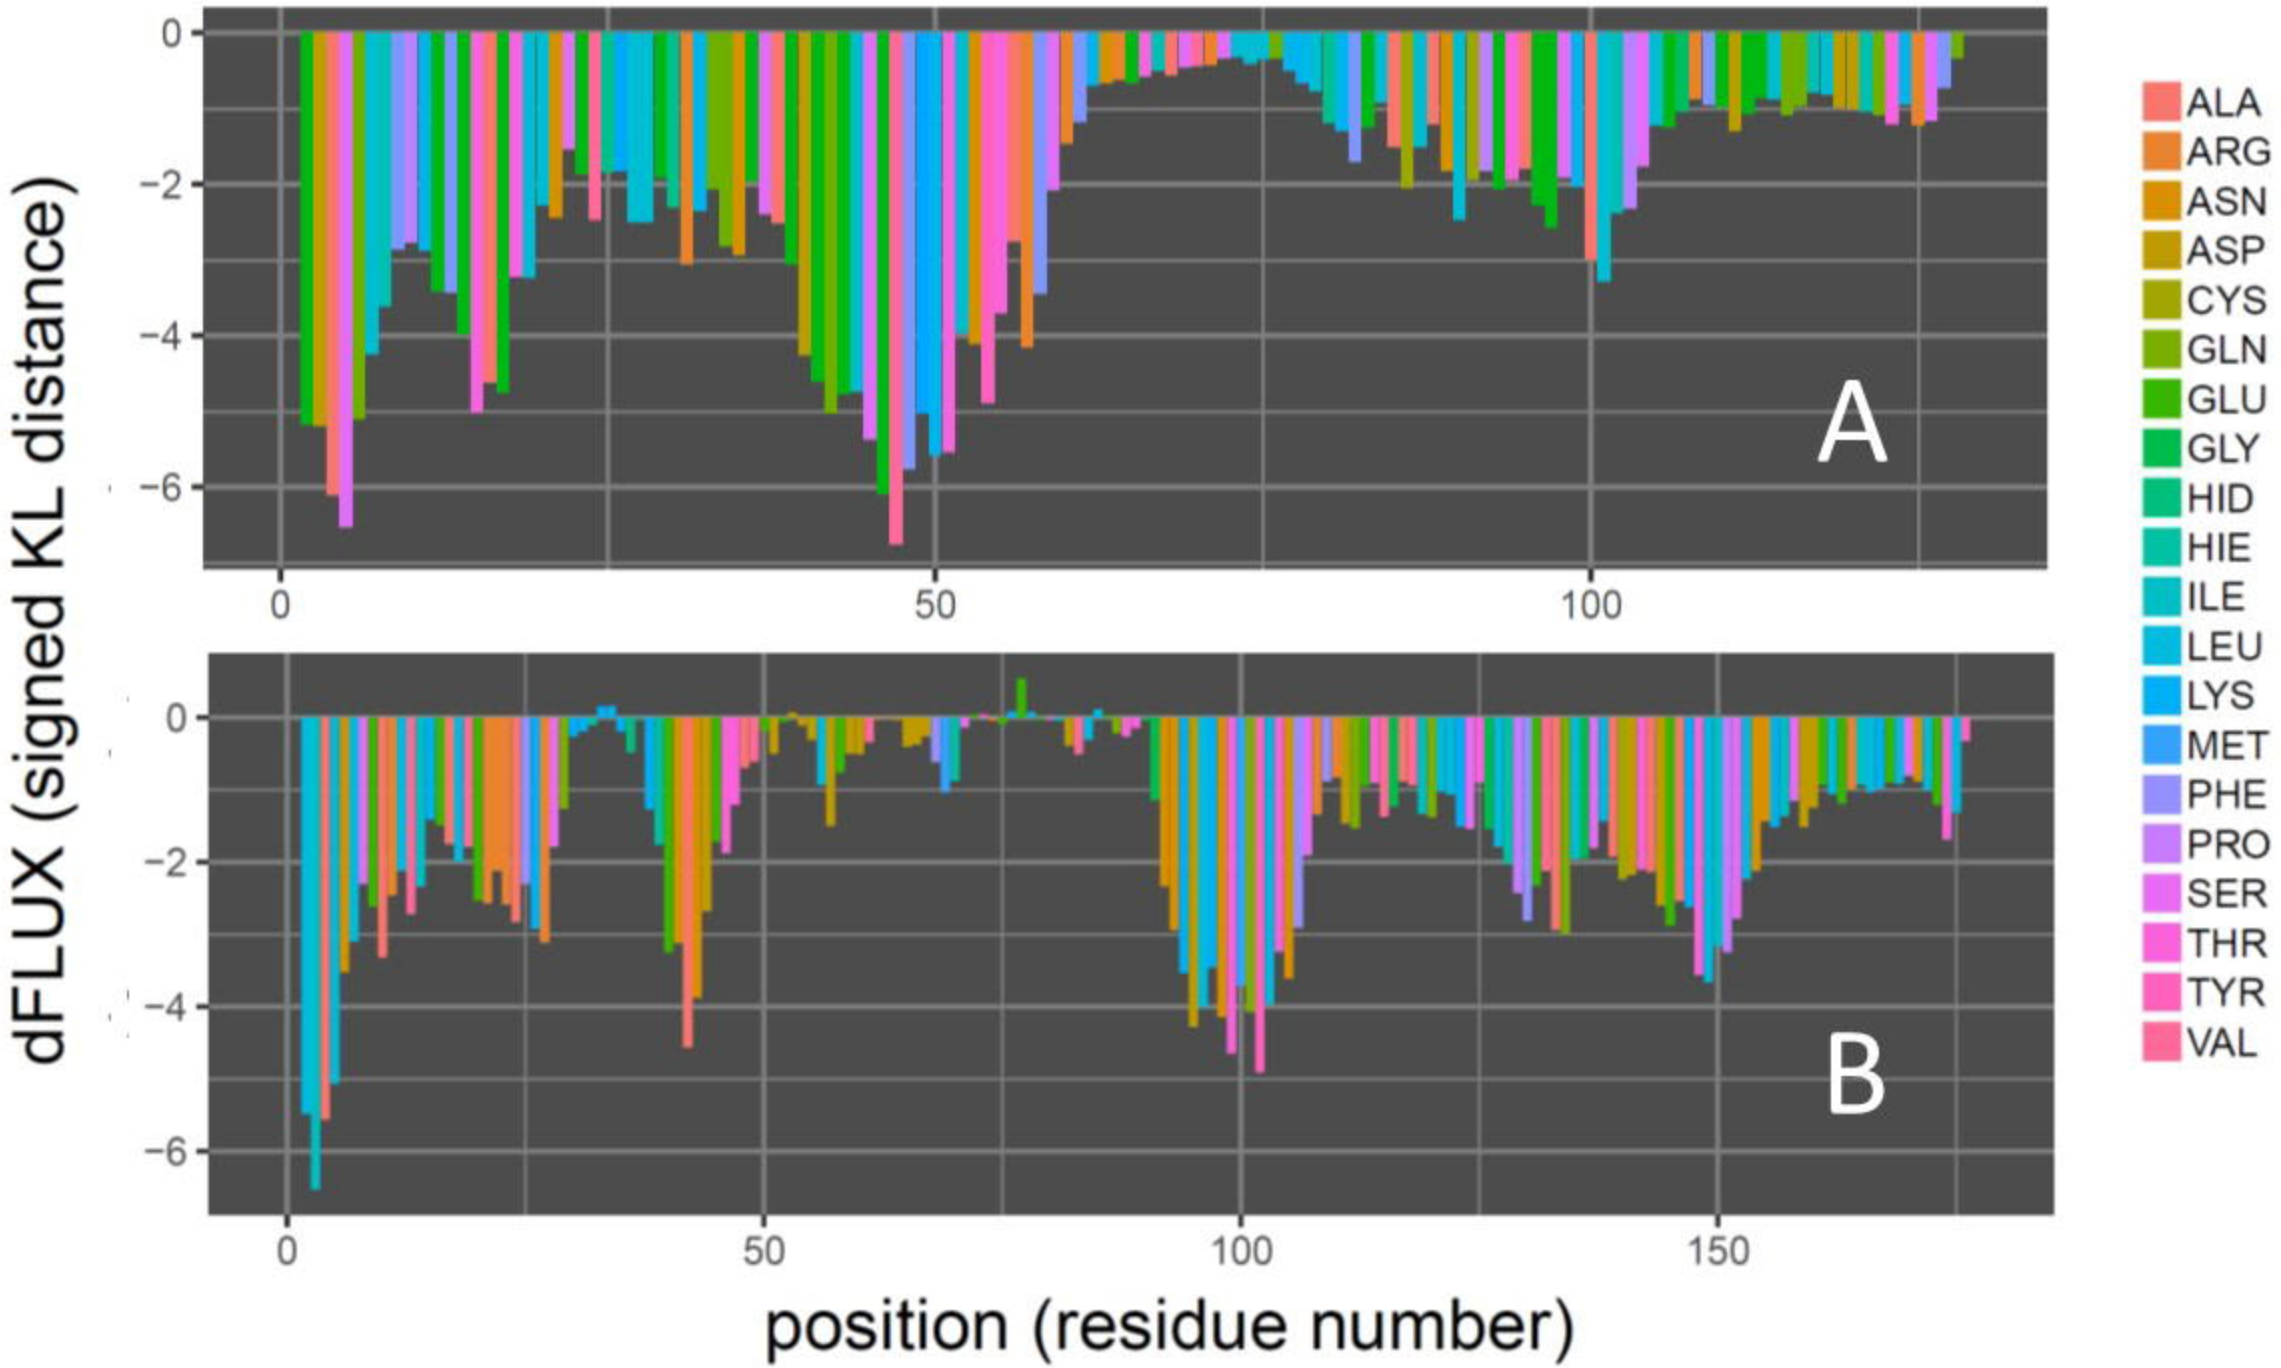
dFLUX (signed KL distance)
0
−2
−4
−6
0
50
100
A
0
−2
−4
−6
0
50
100
150
B
position (residue number)
ALA
ARG
ASN
ASP
CYS
GLN
GLU
GLY
HID
HIE
ILE
LEU
LYS
MET
PHE
PRO
SER
THR
TYR
VAL

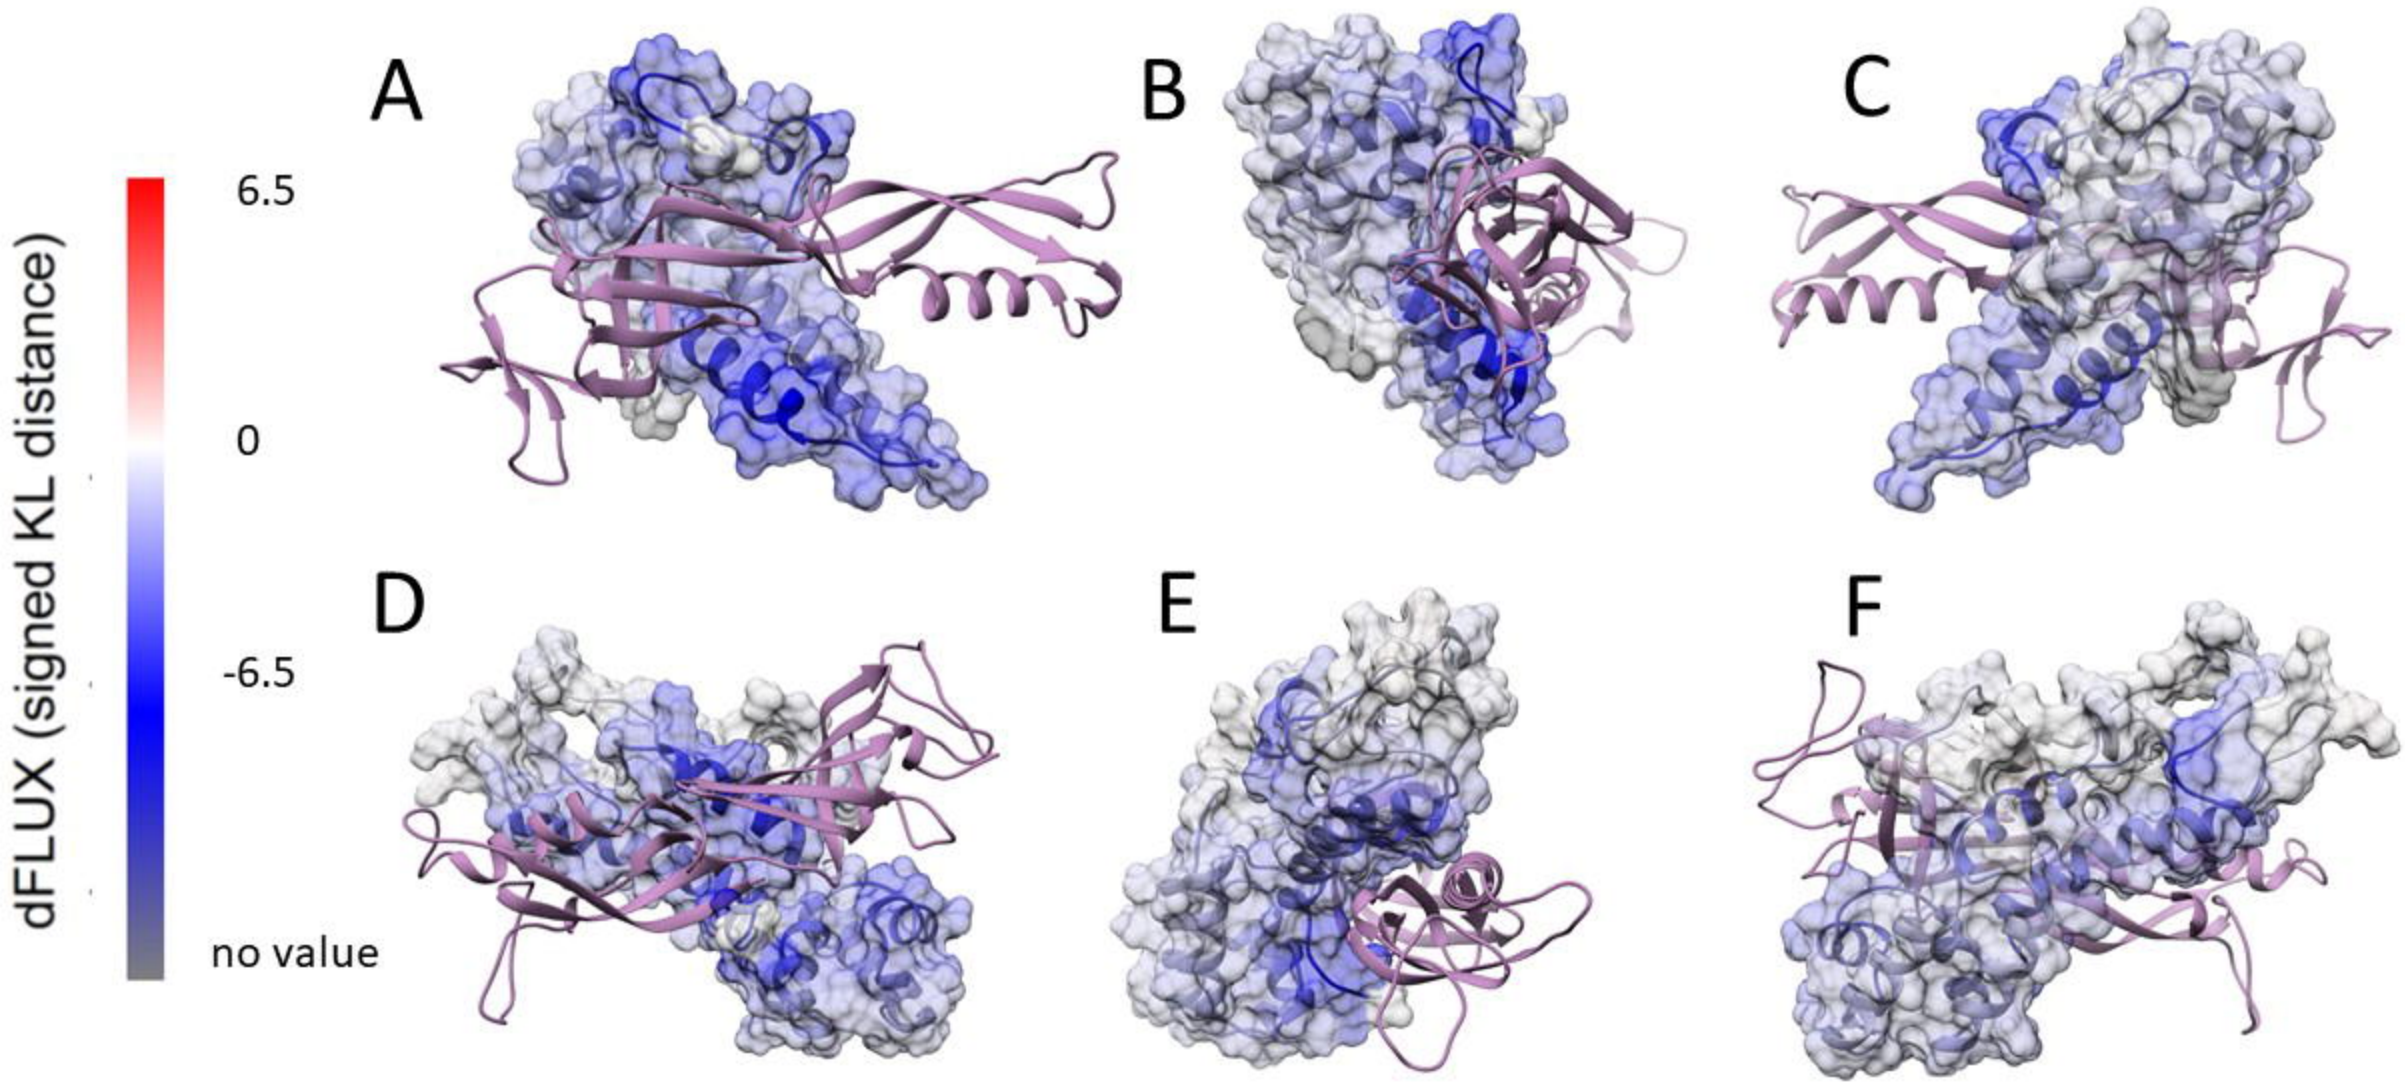
dFLUX (signed KL distance)
6.5
0
-6.5
no value
A
B
C
D
E
F

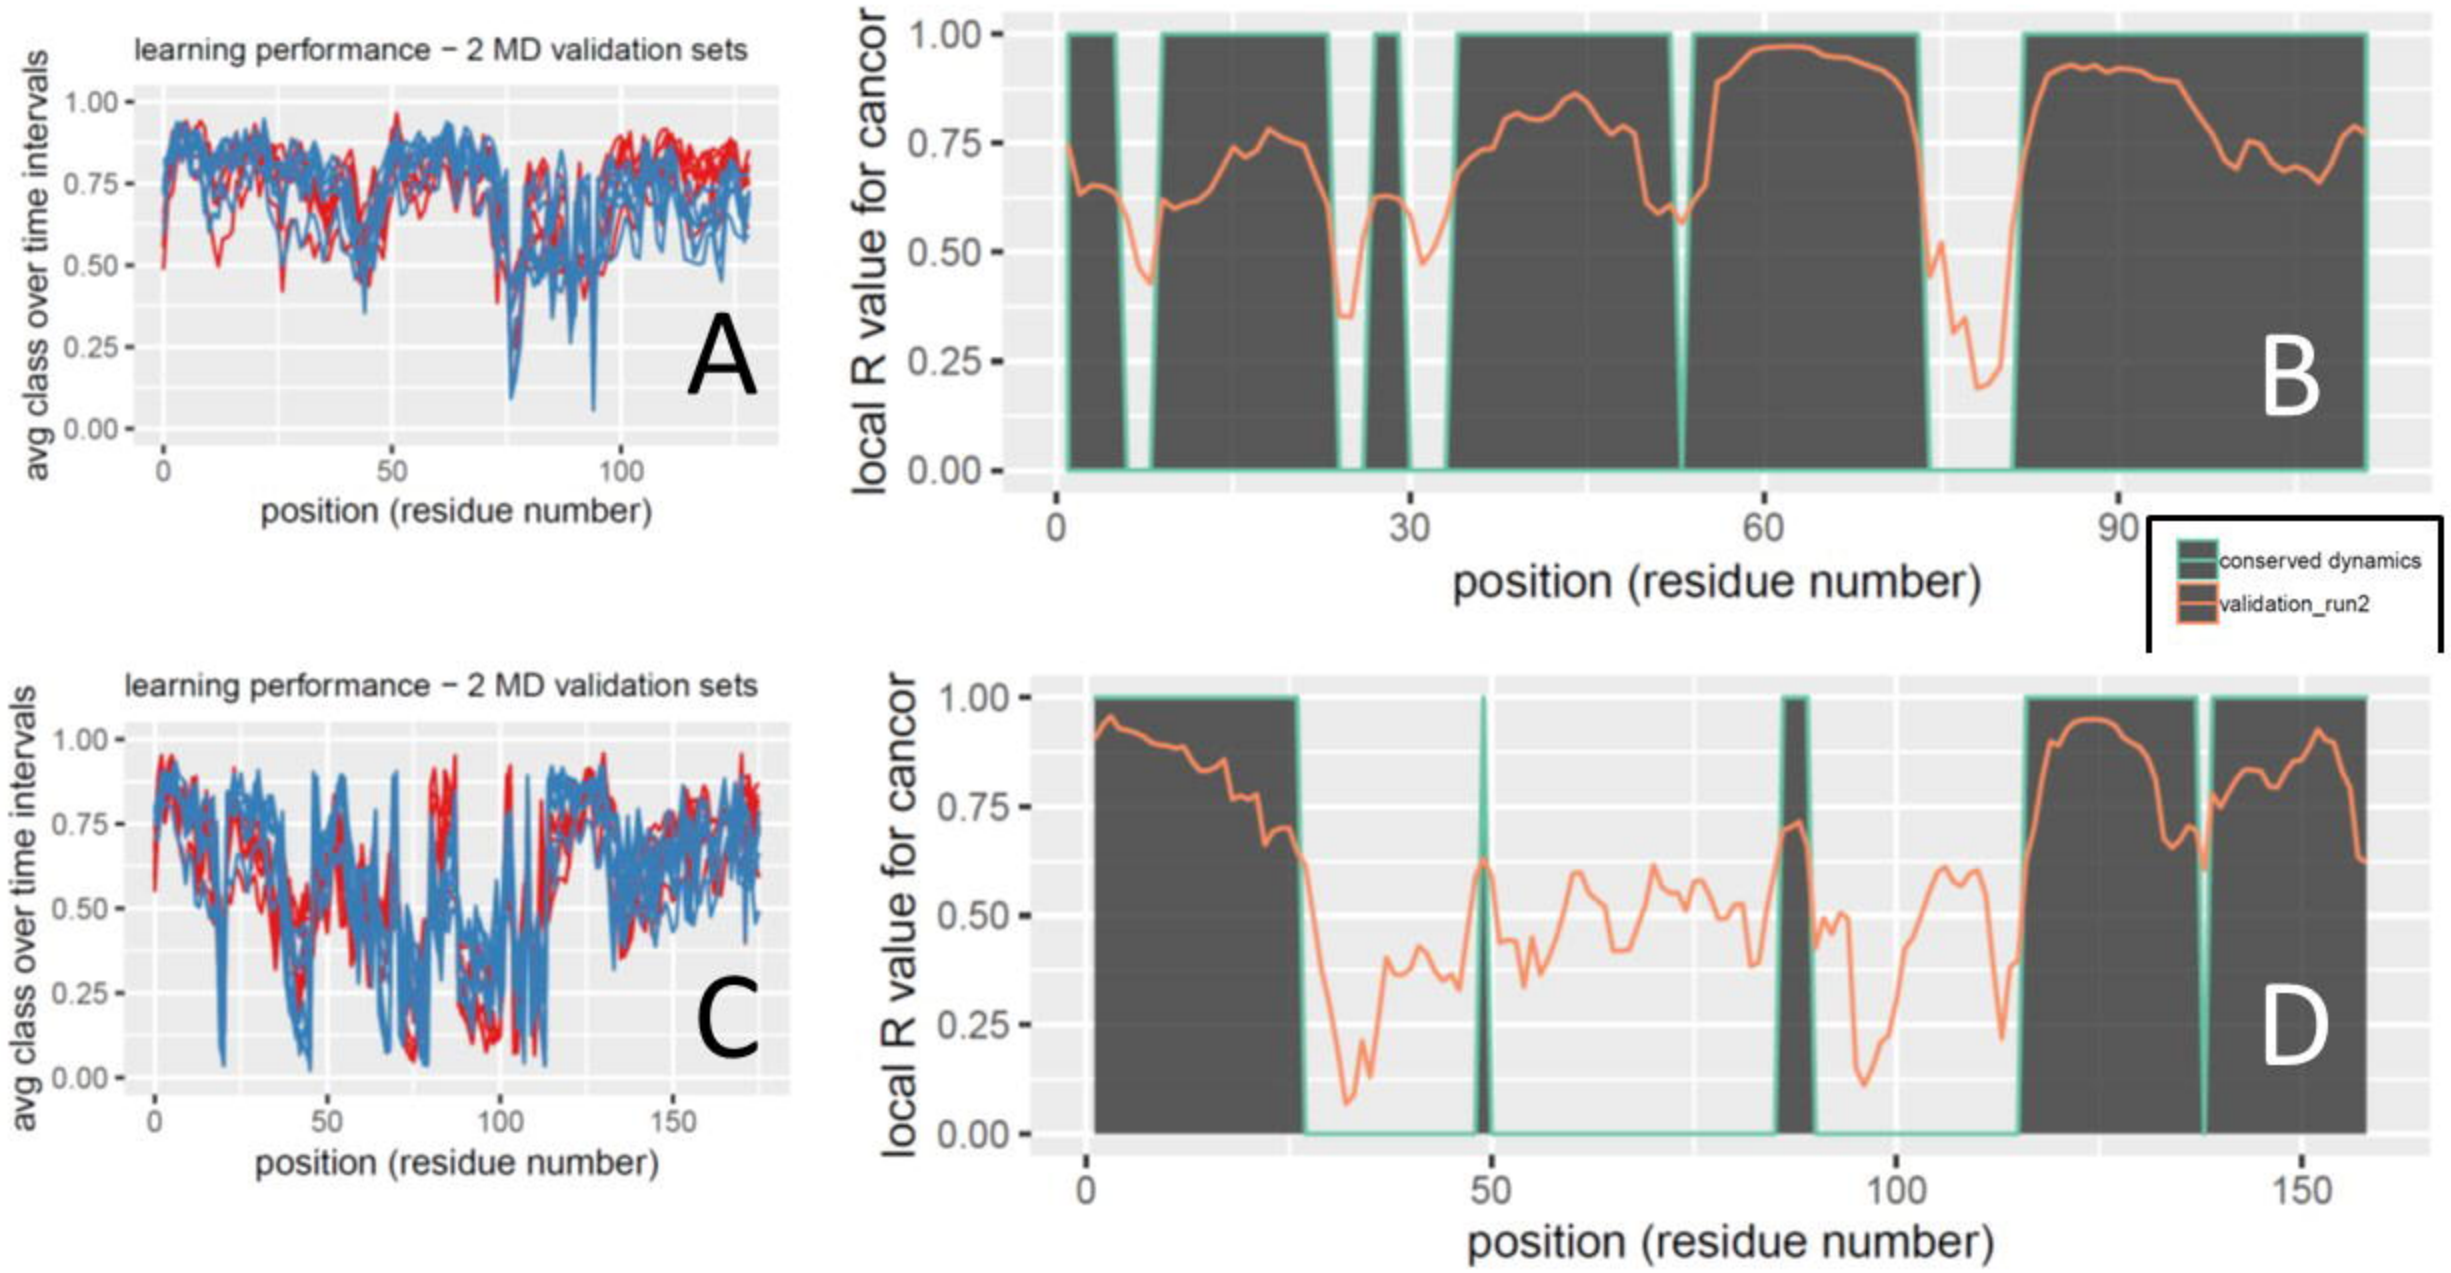
learning performance – 2 MD validation sets
avg class over time intervals
1.00
0.75
0.50
0.25
0.00
0
50
100
position (residue number)
A
local R value for cancor
1.00
0.75
0.50
0.25
0.00
0
30
60
90
position (residue number)
B
conserved dynamics
validation_run2
learning performance – 2 MD validation sets
avg class over time intervals
1.00
0.75
0.50
0.25
0.00
0
50
100
150
position (residue number)
C
local R value for cancor
1.00
0.75
0.50
0.25
0.00
0
50
100
150
position (residue number)
D

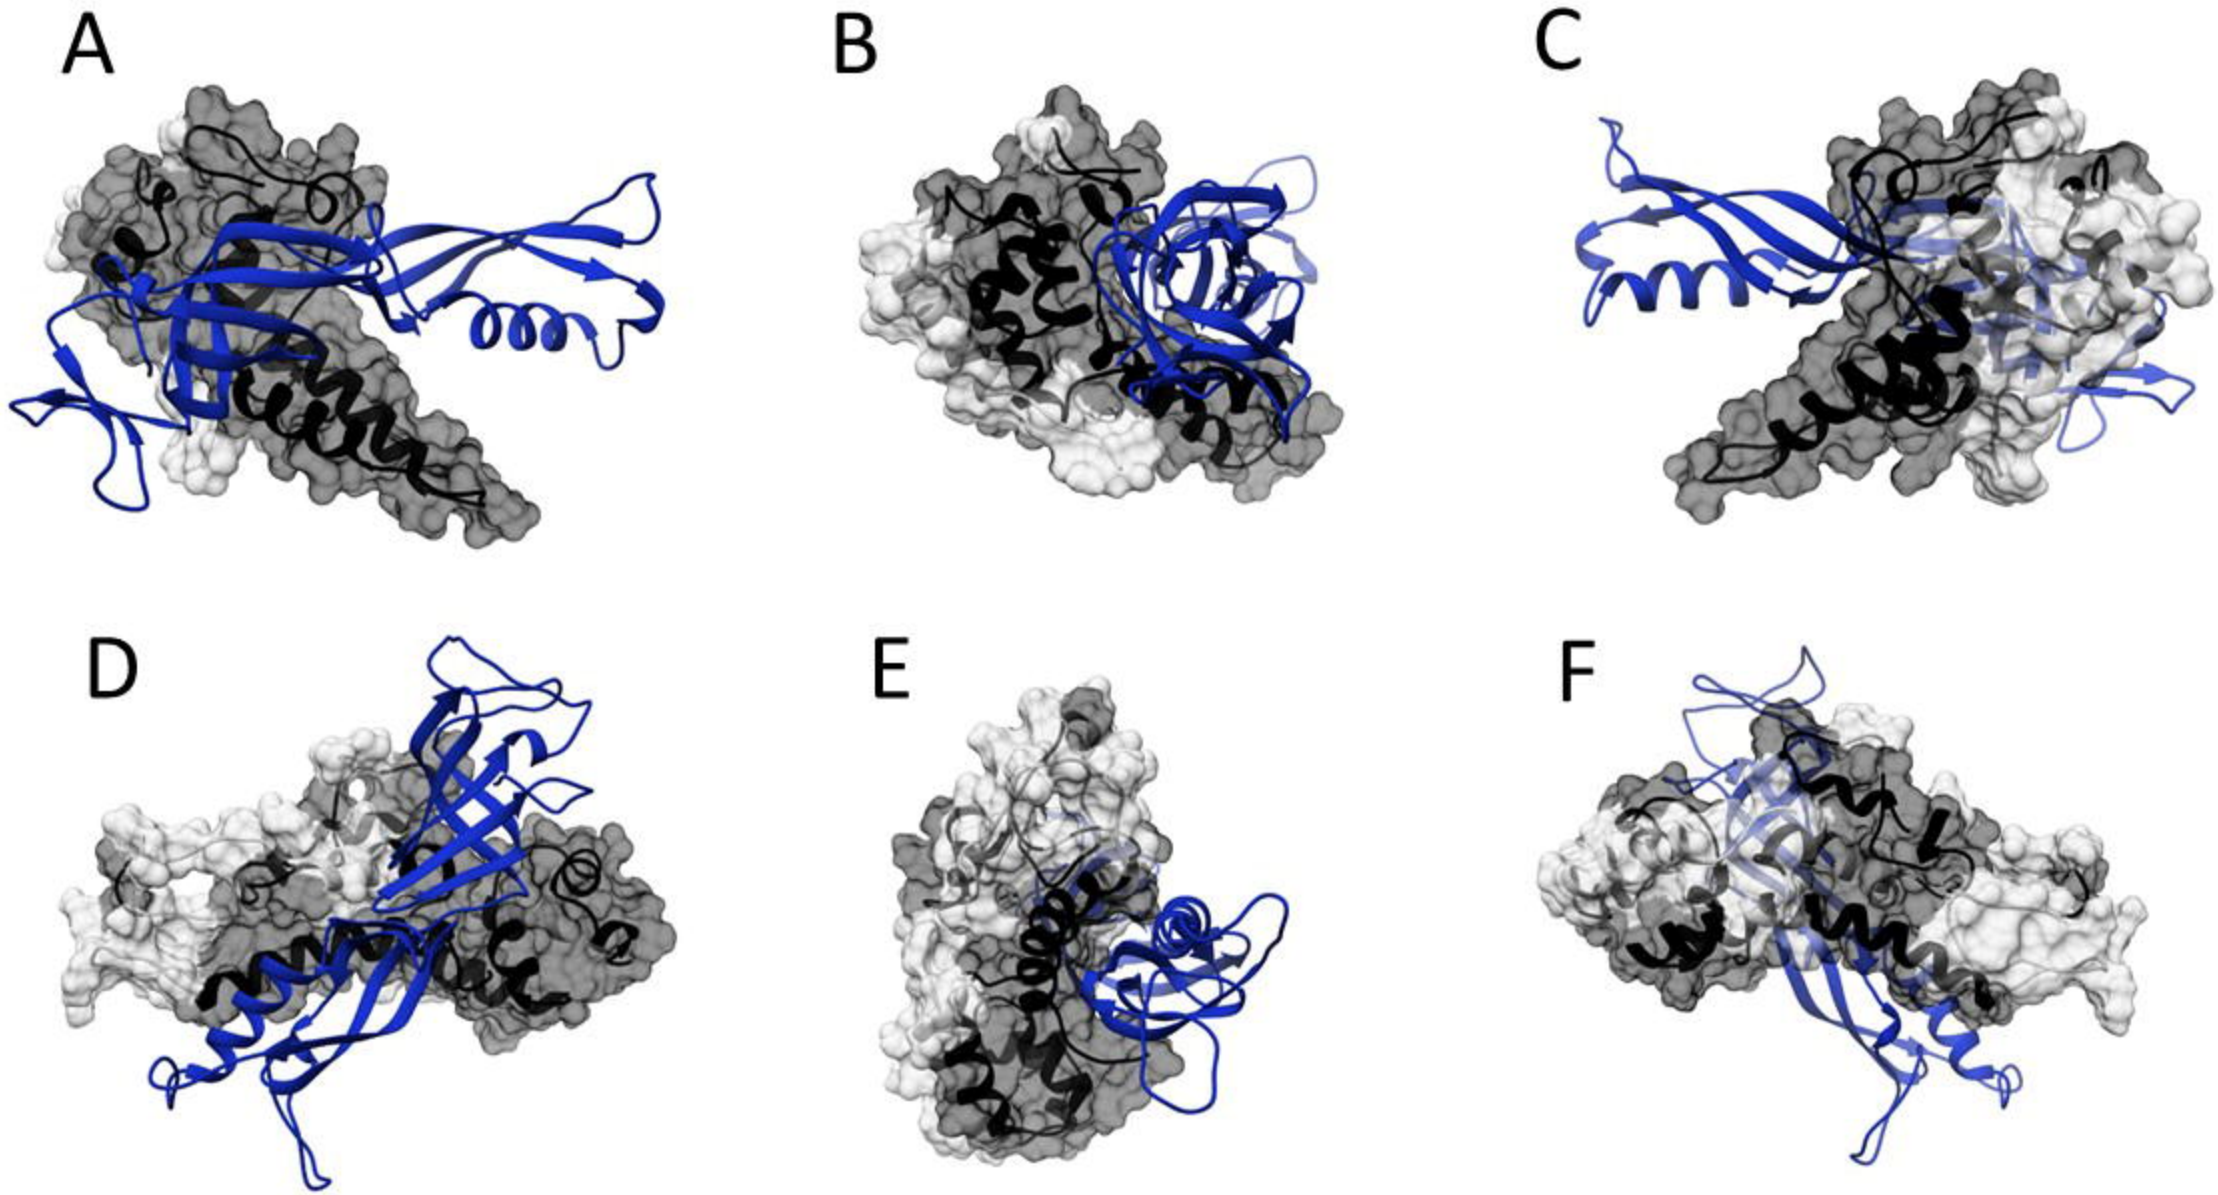
A
B
C
D
E
F

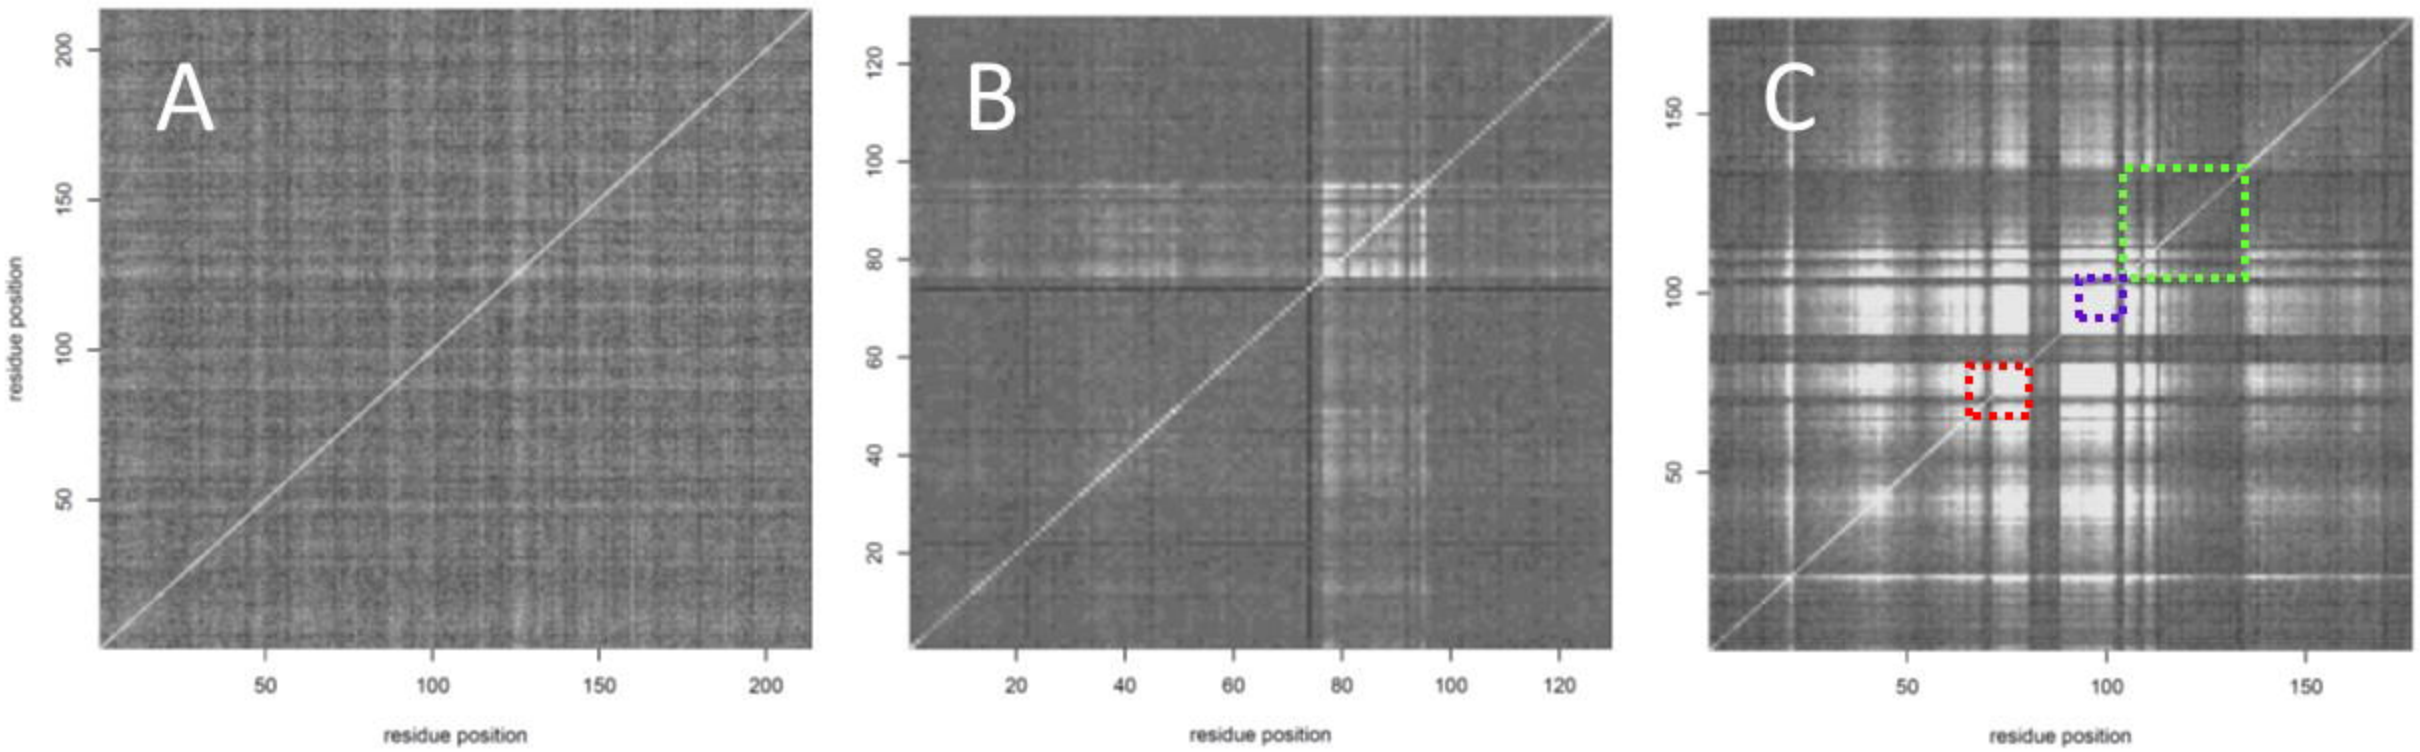
A
B
C
residue position
residue position
residue position
residue position
residue position
residue position

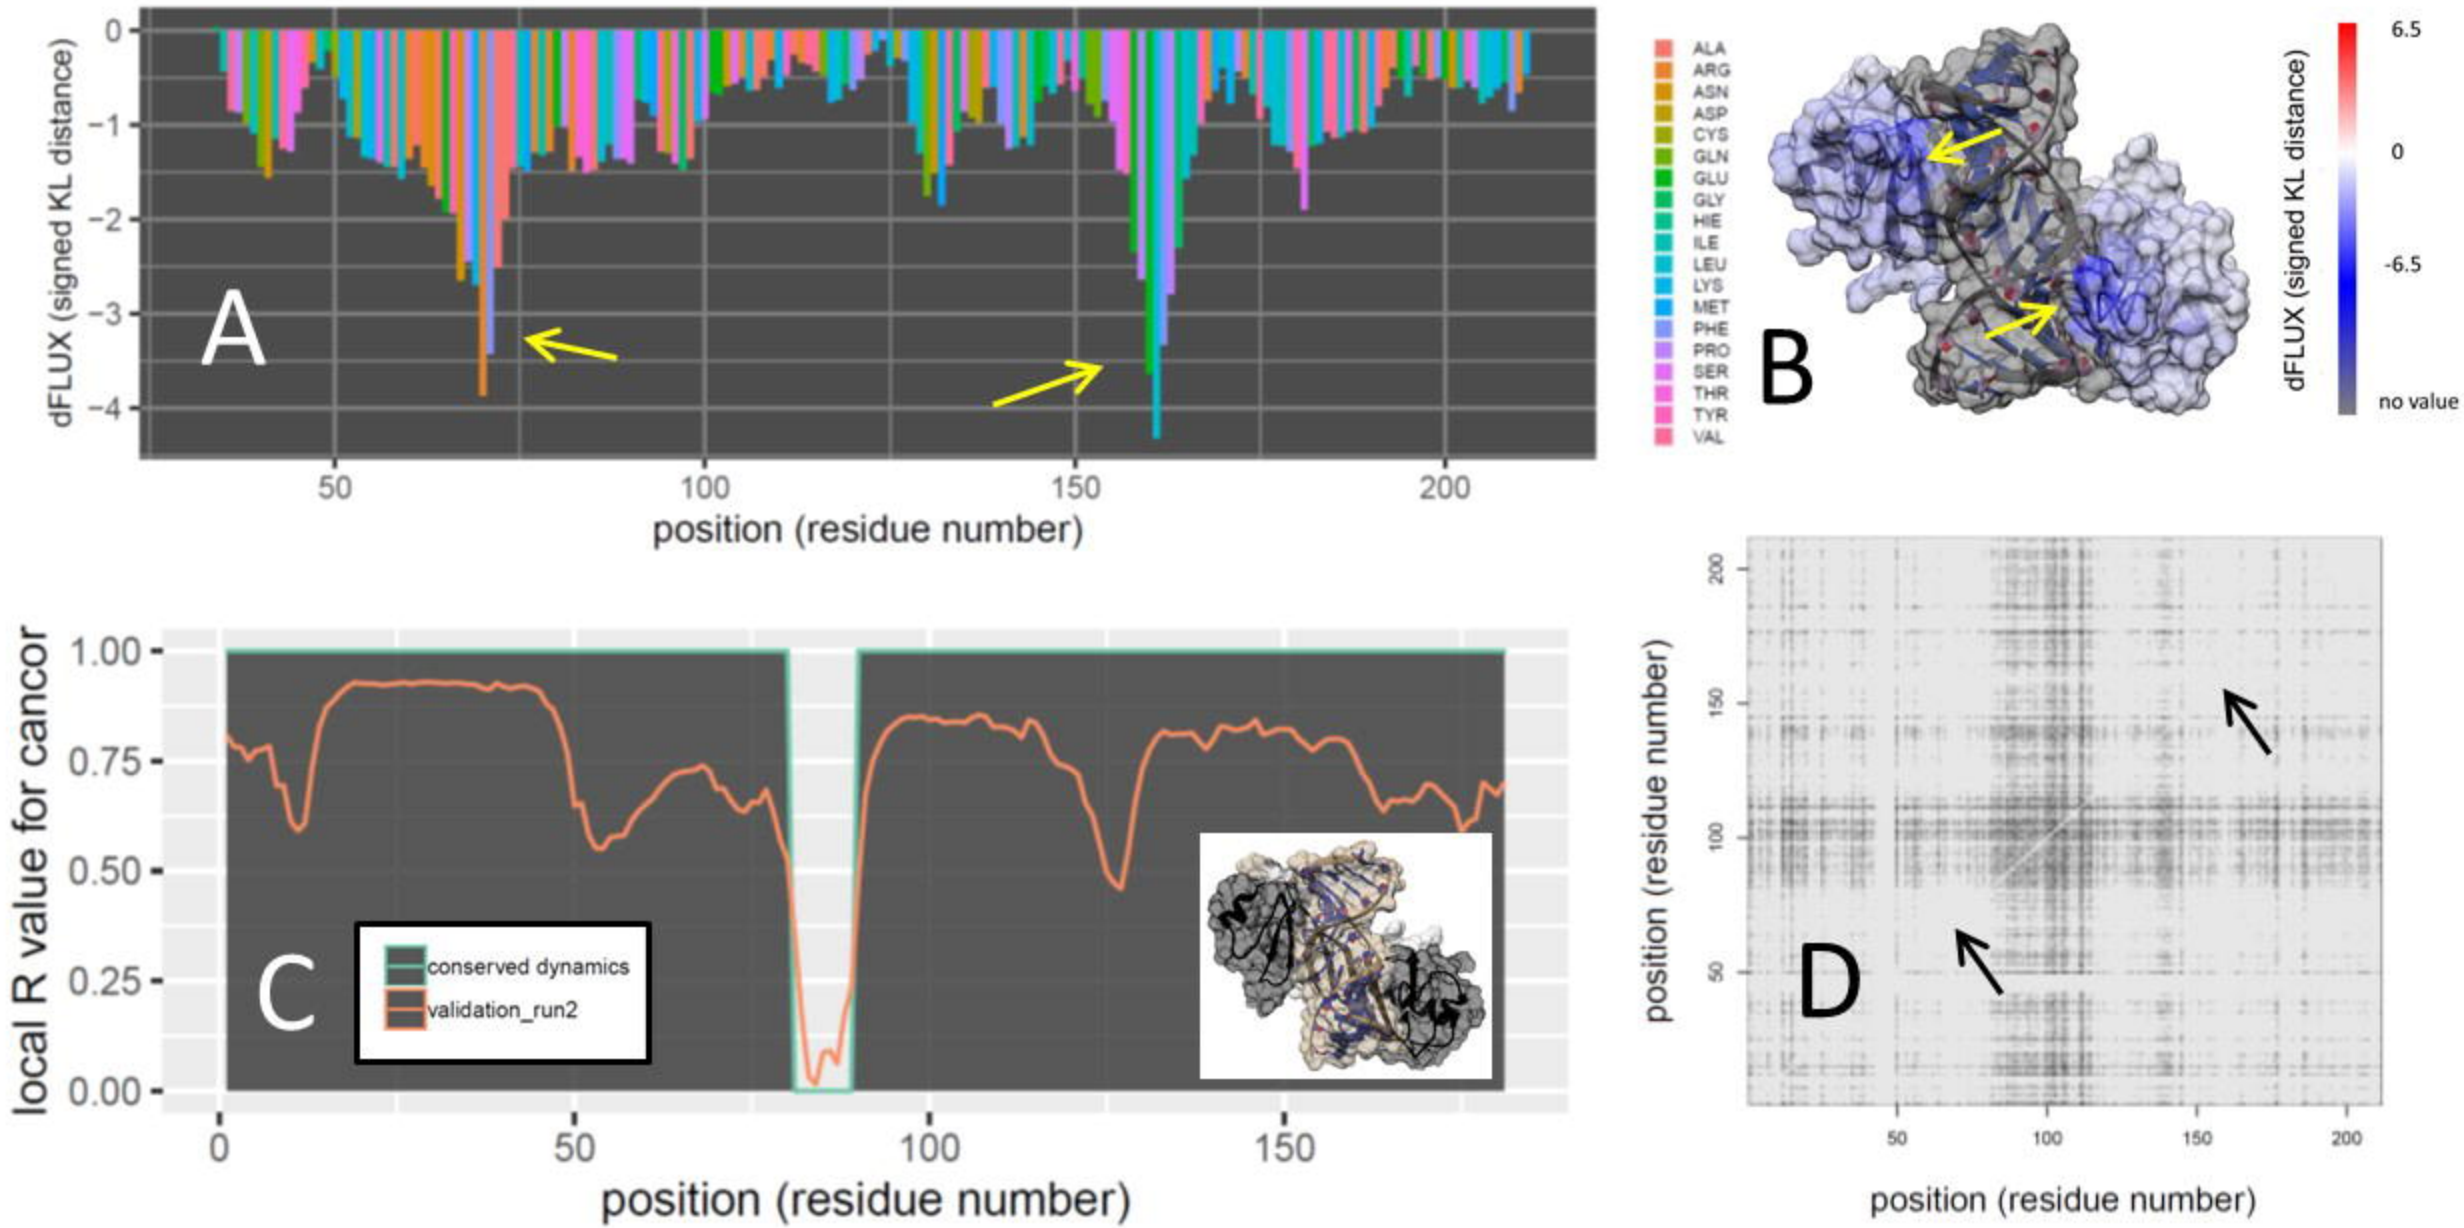
A
dFLUX (signed KL distance)
0
−1
−2
−3
−4
50
100
150
200
position (residue number)
B
ALA
ARG
ASN
ASP
CYS
GLN
GLU
GLY
HIE
ILE
LEU
LYS
MET
PHE
PRO
SER
THR
TYR
VAL
dFLUX (signed KL distance)
6.5
0
-6.5
no value
C
local R value for cancor
1.00
0.75
0.50
0.25
0.00
conserved dynamics
validation_run2
0
50
100
150
position (residue number)
D
position (residue number)
50
100
150
200
position (residue number)