Citation for published version:

Jimenez-Roldan, JE, Freedman, RB, Römer, RA & Wells, SA 2012, 'Rapid simulation of protein motion: merging flexibility, rigidity and normal mode analyses', *Physical Biology*, vol. 9, no. 1, 016008.
<https://doi.org/10.1088/1478-3975/9/1/016008>

DOI:

[10.1088/1478-3975/9/1/016008](https://doi.org/10.1088/1478-3975/9/1/016008)

Publication date:

2012

Document Version

Peer reviewed version

[Link to publication](#)

© IOP

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Rapid simulation of protein motion: merging flexibility, rigidity and normal mode analyses

J. E. Jimenez-Roldan^{1,2}, R. B. Freedman², R. A. Römer¹, S. A. Wells¹

¹ Department of Physics and Centre for Scientific Computing, University of Warwick, Coventry, CV4 7AL, UK

² School of Life Sciences, University of Warwick, Coventry, CV4 7AL, UK

E-mail: j.e.jimenez@warwick.ac.uk

Abstract. Protein function frequently involves conformational changes with large amplitude on time scales which are difficult and computationally expensive to access using molecular dynamics. In this paper we report on the combination of three computationally inexpensive simulation methods — normal mode analysis using the elastic network model, rigidity analysis using the pebble game algorithm, and geometric simulation of protein motion — to explore conformational change along normal mode eigenvectors. Using a combination of ELNEMO and FIRST/FRODA software, large-amplitude motions in proteins with hundreds or thousands of residues can be rapidly explored within minutes using desktop computing resources. We apply the method to a representative set of 6 proteins covering a range of sizes and structural characteristics and show that the method identifies specific types of motion in each case and determines their amplitude limits.

Revision : 1.72, compiled 5 August 2011

Protein Rigidity, protein flexibility, normal mode analysis, conformational change, domain motion, rapid simulation, coarse-grained methods, multi-scale methods.

1. Introduction

Protein conformational changes and dynamic behaviour are fundamental for processes such as catalysis, regulation, and substrate recognition. The time scales of motions involved in enzyme function span multiple orders of magnitude, from the picosecond timescale of local side groups rotations to the milli-/microsecond timescale of the motion of entire domains [1]. Empirical-potential molecular dynamics (MD) has proved to be a valuable tool for investigating molecular motions, but specialized expertise, large-scale computing resources and weeks or months of compute time are required to explore protein motion on simulation time scales greater than tens of nanoseconds [2]. There is clearly a need for methods that permit exploration of possible large conformational motions of proteins in a rational, albeit somewhat simplified, fashion with minimal computational resources. Such explorations allow for the generation of hypotheses about

large conformational motion and protein function that can then be investigated with detailed MD simulations and by experimental methods such as FRET [3].

The computational cost of simulations can be reduced by coarse-graining (CG) — averaging out atomic degrees of freedom so as to represent groups of atoms by a single site [4, 5] — and/or by simplifying the intersite interactions. Different levels of simplification can then be combined in multi-scale methods [6, 7]. Here, we shall consider three such methods in particular. Pebble-game rigidity analysis, implemented in FIRST [8], provides valuable information on the distribution of rigid and flexible regions in a structure [9]. Geometric simulation in the FRODA algorithm [10] uses rigidity information and explores flexible motion [11, 12]. Normal mode analysis of a coarse-grained elastic network model (ENM), implemented in ELNEMO [13, 14], generates eigenvectors for low-frequency motion which are potential sources of functional motion and conformational change [15–19].

Our approach in this study is to bias the generation of new conformations in FIRST/FRODA along an eigenvector predicted by ELNEMO as a low-frequency mode. The bias directs the motion of the atoms in the direction of the eigenvector while the geometric constraint system maintains rational bonding and sterics and prevents the build-up of distortions that occur in a linear projection. The method is outlined schematically in Figure 1. We apply our method to a set of six proteins of various sizes from 58 to 1605 residues. We find that flexible motion in an all-atom model can be explored to large amplitudes in a few CPU-minutes, until further motion is limited by bonding or steric constraints, representing the calculated limit of motion along that vector.

2. Methods

2.1. Protein selection

We deliberately selected six proteins for analysis that are diverse in function, structural characteristics and size, ranging from 58 to 1605 residues. For each protein, we selected a representative high-resolution structure from the Protein Data Bank (PDB) [20]. The proteins and their PDB codes are listed in Table 1 and their structures are shown in Figure 2 with colour coding according to the results of rigidity analysis.

Bovine pancreatic trypsin inhibitor (BPTI) is a small well-studied protease inhibitor of 58 amino acids, comprising mainly random-coil structure plus two antiparallel β -strands and two short α -helices; the protein has only a small hydrophobic core, but is additionally stabilized by 3 intra-chain disulphide bonds [21, 22]. Mammalian mitochondrial cytochrome-c is a classic electron-transfer protein containing a redox-active haem group bound within a primarily α -helical protein fold. These two were selected as contrasting small proteins.

As medium size proteins we selected α 1-antitrypsin and the core catalytic domain of the motor protein kinesin [23]. The former is a protease inhibitor of the serpin family [24]

which operates via a ‘bait’ mechanism comparable to that of a mouse-trap, involving a very significant conformational change, whereas the latter is a mechanochemical device that transduces the chemical energy of ATP hydrolysis into mechanical work, specifically the depolymerisation of microtubules in the case of this kinI kinesin. Both these proteins comprise an extensive β -sheet core flanked by several α -helices.

Protein disulphide-isomerase (PDI) is a large protein (with more than 500 residues) comprising 4 distinct domains each with a thioredoxin-like fold, connected by two short and one longer linker [25]; the protein has both redox and molecular chaperone activity and intramolecular flexibility is essential for its action in facilitating oxidative folding of secretory proteins [26, 27]. The largest protein selected is an integral membrane protein (a bacterial protein of 1605 residues) that operates as a pentameric ligand-gated ion channel (pLGIC); it comprises an extracellular — mainly β -sheet — domain and a membrane-embedded domain, mainly comprising α -helices which form the lining of the ion-channel; the mechanisms of ion permeation and channel gating are not yet completely understood but it is clear that a conformational change is required for function [28].

2.2. Rigidity analysis and energy cutoff selection

We add the hydrogen atoms absent from the PDB X-ray crystal structures using the software REDUCE [29] and remove alternate conformations and renumber the hydrogen atoms in PYMOL [30]. This produces usable files for FIRST rigidity analysis. For each protein we produce a “rigidity dilution” or rigid cluster decomposition (RCD) [8] plot (displayed in the supplementary Figure S1). The plots show the dependence of the protein rigidity on an energy cutoff parameter, E_{cut} , which determines the set of hydrogen bonds to be included in the rigidity analysis. The tertiary structures with the residues coloured by the rigid clusters they belong to are shown in Figures S2–S7 for each of the selected energy cutoffs.

Previous studies [8] suggested that E_{cut} should be at least -0.1 kcal/mol in order to eliminate a large number of very weak hydrogen bonds, and that a natural choice is near the ‘room temperature’ energy of -0.6 kcal/mol. We have shown recently that this criterion is not sufficient to avoid sensitivity to protein-specific structural variations [9] and have argued that E_{cut} must be chosen on a case-by-case basis so as to test specific hypotheses about the rigidity of a particular structure. For each protein we have selected several energy cutoffs at which to explore flexible motion, as listed in Table 1. A higher cutoff energy increases the number of constraints included in the simulation, and this is expected to restrict protein motion. We have in each case used at least one cutoff at which the protein is largely rigid and at least one lower cutoff at which the protein is largely flexible.

2.3. Normal modes of motion

We obtain the normal modes of motion using the ENM [31] implemented in ELNEMO software [13,14]. This generates, for each protein, a set of eigenvectors and associated eigenvalues. Other implementations of elastic network models are also available, for example the AD-ENM of Zheng et al. [32].

The low-frequency modes are expected to have the largest amplitudes and thus be most significant for large conformational changes. However, the six lowest-frequency modes (modes 1 to 6) are trivial combinations of rigid-body translations and rotations of the entire protein. For illustration, here we consider the *five* lowest-frequency non-trivial modes, that is modes 7 to 11 for each protein. We will denote these modes as $m_7, m_8, \dots, m_{11}$.

The mode eigenvectors are predicted on the basis of a single protein conformation. The amplitude to which a mode can be projected may be limited by bonding and/or steric constraints that are not evident in the input structure or fully captured by the ENM. A linear projection of all the residues in the protein along a mode eigenvector introduces unphysical distortions of the interatomic bonding. Typically, to avoid this and project a mode to finite amplitude requires one or more cycles of a combined method; the mode is projected linearly until distortions become evident and the resulting structure is relaxed using constrained MD/molecular mechanics [33–35]. We explore an alternative method for projection of modes to large amplitudes, using rigidity analysis and geometric simulation.

2.4. FRODA mobility simulation

Geometric simulation, implemented in the FRODA module within FIRST [10,36], explores the flexible motion available to a protein with a given pattern of rigidity and flexibility. New conformations are generated by applying a small random perturbation to all atomic positions; FRODA then reapplies bonding and steric constraints to produce an acceptable new conformation. Motion can be biased by including a directed component to the perturbation. The capability to use a mode eigenvector as a bias was implemented in FIRST/FRODA by one of us (SAW) and has been briefly reported previously [37]. The combination of ELNEMO and FIRST/FRODA, illustrated schematically in Figure 1, is described in detail in supplementary material (section S1).

Since the displacement from one conformation to the next is small, we record only every 100th conformation and continue the run for typically several thousand conformations. The run is considered complete when no further projection along the mode eigenvector is possible (due to steric clashes or bonding constraints) which manifests itself in slow generation of new conformations and poor reproducibility in the results of independent runs. We have performed FRODA mobility simulation for each protein at several selected values of E_{cut} , see section 2.2.

During conformation generation we track the fitted RMSD between α carbons of the initial and current conformation. This measure is discussed in more detail in

supplementary material (section S2). To project a mode to an amplitude of several Å RMSD typically takes a few CPU-minutes on a single processor. We carry out five parallel simulations for each structure, mode and direction of motion and monitor the evolution of RMSD during each run, as illustrated in Figures 3-6.

3. Results

3.1. Conformer generation with mode bias

The output of the mobility simulations for BPTI (1BP1) is summarized in Figures 3a–c. The evolution of RMSD for each of $m_7, \dots, m_{11}$, during runs at two selected values of E_{cut} , is represented in Figures 3b and 3c respectively. In all cases, we observe an initial phase in which the RMSD increases almost linearly, as the protein explores the mode direction without encountering significant steric or bonding constraints on the motion. During this phase, generation of new conformations in FRODA is very rapid and the RMSDs from different runs are very similar to each other. The RMSD then displays an inflection, ceasing to rise linearly, and approaching an asymptote; this indicates that steric clashes and bonding constraints (such as hydrophobic tethers) are preventing further exploration along the mode direction. The asymptote is thus an amplitude limit on the mode. In this phase the generation of new conformations in FRODA becomes slower as the fitting algorithm has increasing difficulty finding a valid conformation, and the RMSDs achieved by different runs differ. In the regime of slow conformation generation, the mode bias is forcing the structure into a regime of steric clashes and/or of bonding constraint limits, for example when residues connected by a hydrophobic tether are being pushed apart. This regime cannot correspond to the low-frequency flexible motion which we wish to explore. Our presentation of RMSD data for larger proteins is therefore truncated once this “jamming” starts. For most of our proteins, 2500 conformations is sufficient to cover the regime of rapid conformation generation. For our largest protein, pLGIC (2VL0), we find that 1000 conformations was sufficient.

3.2. Small loop motion

In Figure 3a we show an ensemble of structures for BPTI generated by exploring the lowest-frequency nontrivial mode, m_7 , and in Figures 3b,c we show the RMSDs achieved for the five lowest-frequency nontrivial modes. The amplitudes of flexible motion for BPTI at the cutoffs -0.2 kcal/mol and -2.2 kcal/mol reach an asymptote at RMSD values around 2 to 4 Å. Considerable asymmetry, a factor of 2 difference in the achievable RMSD, is observable in some modes between the two possible directions of motion. Let us emphasize here that in general the amplitudes of flexible motion are not available from either the elastic network model or the rigidity analysis without simulation of flexible motion.

BPTI (1BPI) is a small protein with some relatively flexible loop regions. For contrast, we now examine a compact globular protein, cytochrome-c (1HRC). RMSD

results are shown in Figure 3e,f truncated once jamming had begun. Although this protein is twice as large as BPTI in terms of residues, its capacity for flexible motion is visibly more limited, with amplitudes below $2\text{\AA}$ in all modes. This result is important in validating our method of projecting modes to large amplitude; if geometric simulation were capable of reaching unphysically large amplitudes, the method would lose its value.

3.3. Large loop motion

RMSDs for low-frequency modes of internal kinesin motor domain protein (1RY6) are shown in Figure 4b, c for two different energy cutoffs. The amplitudes achievable for these modes differ little between the different energy cutoffs; motion occurs principally in a flexible loop region around residues 37–46. An exploration of m_7 is presented in Figure 4a for an energy cutoff of -1.1 kcal/mol, clearly showing the loop motion. We find that the combination of the mode bias and the bonding constraints naturally causes the large flexible loop to follow a curved trajectory.

As shown in Figures 4e, f, $\alpha 1$ -antitrypsin (1QLP) displays several low-frequency modes which easily explore amplitudes of up to $2\text{--}2.5\text{\AA}$ RMSD depending on the rigidity cutoff. The motion shown in Figure 4d again involves the easy motion of large flexible loops with respect to the relatively rigid β -sheet core of the protein.

3.4. Domain motion

Protein disulphide isomerase (2B5E) is an interesting case for protein mobility. The protein consists of four domains (a–b–b'–a') connected by flexible linkers. Biological evidence indicates that conformational flexibility is vital to the function of the enzyme [25]. Rigidity analysis immediately brings out the flexibility of the molecule (see Figure S1e). Even at very high E_{cut} values ($E_{\text{cut}} = -0.015$ kcal/mol), the rigidity analysis reveals the domain organisation of the protein with each domain corresponding to a distinct rigid cluster flanked by flexible linkers. The RMSD achievable by low-frequency flexible motion therefore does not depend significantly on the energy cutoff; motion is slightly limited at the weakest cutoff (Figure 5b), but at other cutoffs the achievable amplitudes are essentially the same (Figure 5c and d). Close examination of the conformation generation in FIRST/FRODA indicates that the amplitudes are limited eventually as further motion along the mode would over-extend covalent and hydrophobic-tether constraints.

The inter-domain nature of the flexible motion is detailed in Figure 5a which shows structures at the amplitude limits of m_7 in the positive and negative directions. The structures are aligned on the b–b' domains, bringing out the motion of the a domain and particularly of the a' domain. The CPU time required to project this protein with more than 500 residues along m_7 to an amplitude of more than $20\text{\AA}$ is less than 15 minutes.

The largest protein we have investigated is pLGIC (2VL0). Its pentameric structure includes a transmembrane domain composed of α -helices and an extracellular domain consisting largely of β -sheets. The major rigidity transition in the protein identified by

FIRST occurs between a cutoff of -0.4 kcal/mol, when almost the entire structure forms a single rigid cluster, and -0.5 kcal/mol, when the five backbone sections linking the two major domains have become flexible and the many α -helices in the transmembrane are mutually flexible. At the lower energy cutoff it is possible for the two domains to move relative to each other, and the transmembrane helices are also capable of relative motion. The increased RMSD possible at the lower E_{cut} is visible in Figure 6.

The flexible motion at the higher cutoff, with the protein largely rigid, involves only the motion of a few flexible loops. The motion at the lower cutoff is far more biologically interesting. The lowest-frequency non-trivial mode, m_7 , involves a counter-rotation of the transmembrane and extracellular domains, including a change in the relative tilt of transmembrane helices lining the ion channel. This flexible motion is shown in Figure 6a, b. The CPU time to project this large protein with more than 1600 residues along m_7 to its amplitude limit is less than twenty minutes.

4. Discussion

4.1. Extensive RMSD as a measure of total flexible motion

The RMSD values displayed in Figures 3–6 range from $1.5\text{\AA}$ to $10\text{\AA}$. However, the character of the flexible motion does not seem well reflected by the RMSD values. A small protein without large conformational change, BPTI shows an RMSD of up to $3.5\text{\AA}$ in its small loop motion (Figure 3); the channel protein pLGIC is thirty times larger by residue count and shows a substantial domain motion, yet shows maximum RMSD values of around $2.5\text{\AA}$. So, although RMSD is a good measure for comparing two similar structures, it does not necessarily capture the scale of motion between two different structures.

In order to quantify the total extent of flexible motion, we introduce an *extensive* RMSD measure by multiplying the RMSD (which describes the average displacement of atoms) by the number of residues in the protein. Figure 7 shows these xRMSD values for all the selected proteins moving along $m_7^\ddagger$. The three categories of motion which we have discussed — small loop motion, large loop motion and domain motion — become clearly visible in xRMSD. The xRMSD results for BPTI and cytochrome-c closely resemble each other even though cytochrome-c is almost double the size of BPTI. Similarly, the kinesin protein and the $\alpha 1$ -antitrypsin display similar xRMSD behaviour to each other in their large loop motion. PDI and the pLGIC likewise have similar xRMSD behaviour reflecting their domain motion.

[‡] For the proteins with domain motion, we have chosen values of E_{cut} which correspond to lower flexibility. The observed large variation in xRMSD is therefore taking place despite a restrictive bond network. On the other hand, for the proteins with loop motion we selected energy cutoffs which allow more structural flexibility. In this situation, although larger regions of the protein could become mobile, we still only observe localized loop motion.

4.2. Monitoring the evolution of normal modes

It is implicit in normal-mode analysis of protein conformational change that a mode eigenvector should be a valid direction for motion over some non-zero amplitude. For example, Krebs *et al.* [38] have surveyed a large number of known conformational changes, using paired crystal structures, comparing the vector describing the observed conformational change to the low-frequency elastic network mode eigenvectors using a dot product. In many cases the observed change had a large dot product (> 0.5) with only one or two normal modes.

In each of our simulations we use an *initial* normal mode, $m_j^{(i)}$, as a bias throughout the simulation. We calculate a new set of *current* normal modes, $m_j^{(c)}$, for each newly generated conformation. We compute the dot product of the bias vector, $m_j^{(i)}$, with $m_j^{(c)}$, that is, the current normal mode with the same mode number j , as in $m_j^{(i)} \cdot m_j^{(c)}$. Graphs of these dot products are shown for all protein structures investigated here in the supplementary Figures S8–S13.

We find that three main classes of dot product behaviour emerge, shown schematically in Figure 8a. In motif 1, the dot product remains close to 1 throughout the simulation, indicating the initial and current modes remain very similar. In motif 2, there is a gradual decline in the dot product, of quadratic or cosine character; this suggests a gradual rotation of $m_j^{(c)}$ relative to $m_j^{(i)}$. Perhaps the most interesting case is motif 3, in which the dot product $m_j^{(i)} \cdot m_j^{(c)}$ collapses rapidly; this can occur at any point in the simulation, even if the RMSD between the initial and current conformations is small. Examples of these motifs can be observed for all our proteins. We find, e.g. a motif 1 behavior in mode m_7 for PDI (cp. Figure S12b, c, d), a motif 2 behaviour in m_7 for kinesin (cp. Figure S11) and motif 3 character in m_7 for antitrypsin (cp. Figure S10) as well as in BPTI (cp. Figure S8). Similar agreement with all motifs can be found for higher modes, although, as a general tendency, the smooth motifs 1 and 2 become gradually less visible and the more rapid changes exemplified by motif 3 more pronounced.

These sudden collapses do not indicate that the initial normal mode eigenvector has ceased to be a valid direction along which flexible motion is possible. Rather, the eigenvector has ceased to represent a single pure mode; $m_j^{(i)}$ now has significant overlap with multiple other modes $m_k^{(c)}$. This *mode mixing* is illustrated, e.g., in Figure 8b for projection of cytochrome-c (1HRC) along $m_{11}^{(i)}$ at a cutoff energy $E_{\text{cut}} = -1.2$ kcal/mol. After projection in the positive direction over about $0.25\text{\AA}$ RMSD, $m_{11}^{(i)}$ has a large overlap with $m_{11}^{(c)}$ and also with $m_{10}^{(c)}$. After a similar projection in the negative direction, $m_{11}^{(i)}$ has only a small overlap with $m_{11}^{(c)}$, and instead has some overlap with many low-frequency modes $m_k^{(c)}$ such as $k = 8, 9, 10, 12, 13, 14$. Thus, a pure mode calculated on one structure may be a mixture of multiple modes when calculated on a very similar conformation.

These results clarify that while the dot products provide useful additional information about the stability of the initial modes during the simulated motion, they are

not simply correlated with loop or domain motion in contradistinction to the xRMSD.

4.3. Significance of rigidity-analysis energy cutoff

In the case of small loop motion, it is clear that lowering the rigidity-analysis energy cutoff — thus making the structure more flexible — increases the amplitude of flexible motion, as one might expect. The simulation of flexible motion can thus add value to rigidity analysis of protein structures by identifying the constraints that must be eliminated in order for two residues to become independently mobile. In the case of large domain motion, however, the most important criterion appears to be whether the domains are mutually rigid or not.

We can see in the case of PDI (Figure 5) that the amplitude of flexible motion for the lowest-frequency modes is almost unaffected by the choice of the energy cutoff (E_{cut}) provided it is set at a reasonable value of energy cutoff which represents each domain as a number of separate small rigid clusters. This conclusion can also be drawn in the case of the ligand-gated ion channel protein.

5. Conclusions

We have reported a hybrid method to explore protein motion by integrating both rigidity constraints from FIRST and directional information, in the form of low-frequency elastic network mode eigenvectors obtained using ELNEMO, into the geometric simulation method FRODA. The exploration brings out features of the motion that could not be inferred using FIRST or ELNEMO alone. In order to illustrate the method, we have applied it here to a diverse selection of proteins whose flexible motion ranges from small loop motion (BPTI, cytochrome-c) and large loop motion (a kinesin and an antitrypsin) to large motions of entire domains (protein disulphide isomerase and a transmembrane pore protein). Detailed studies of dynamics in relation to function of particular proteins are currently in progress [39, 40]. The combined method can rapidly explore motion to large amplitudes in an all-atom model of the protein structure, maintaining steric exclusion and retaining the covalent and noncovalent bonding interactions present in the original structure. Significant amplitudes of motion are achieved with only CPU-minutes of computational effort even in a pentameric pore protein with more than 1600 residues. The amplitude of motion that can be achieved by flexible loops increases as the rigidity-analysis energy cutoff is lowered. For large-scale motion of domains, the most important criterion is that the energy cutoff should be low enough that different domains do not form a single rigid body. We note that RMSD, a measure of structural similarity, does not properly reflect the scale of flexible motion between different proteins; this is better captured by an extensive measure, xRMSD, which reflects both the size of the protein and the amplitude of its motion. Examination of the behaviour of the elastic network eigenvectors during the motion shows many examples of mode mixing, so that a given vector of motion can change from being a pure mode to a mixed one after quite small

displacements, without losing its character as an “easy” direction for flexible motion.

We believe that the ability to explore large amplitudes of flexible motion in an all-atom model with minimal computational resources will be of great use in biochemistry, structural biology, and biophysics, as a generator of new hypotheses and intuitions about protein structure and function, and as an ally to other simulation methods such as molecular dynamics, Monte Carlo folding simulations [41] and *ab-initio* simulations. Such a study is currently in preparation for PDI [39].

Acknowledgements

We gratefully acknowledge support from the Royal Society through the India-UK International Scientific Seminar scheme. SAW thanks the Leverhulme Trust for support through an Early Career Fellowship. JEJ thankfully acknowledges EPSRC and BBSRC funding support during his PhD.

References

- [1] Henzler-Wildman K and Kern D. Dynamic personalities of proteins. *Nature*, 450(13):964–972, 2007.
- [2] Shaw DE, Maragakis P, Lindorff-Larsen K, Piana S, Dror RO, Eastwood MP, Banks JA, Jumper JM, Salmon JK, Shan Y, and Wriggers W. Atomic-level characterization of the structural dynamics of proteins. *Science*, 330:341–346, 2010.
- [3] Giepmans BNG, Adams SR, Ellison MH, and Tsien RY. Review- the fluorescent toolbox for assessing protein location and function. *Science*, 312:217–224, 2006.
- [4] Russell D, Lasker K, Phillips J, Schneidman-Duhovny D, Velazquez-Muriel JA, and Sali A. The structural dynamics of macromolecular processes. *Curr. Opin. Cell Biol.*, 21:97–108, 2009.
- [5] Clementi C. Coarse-grained models of protein folding: toy models or predictive tools? *Curr Opin Struct Biol*, 18:10–15, 2008.
- [6] Yaliraki SN and Barahona M. Chemistry across scales: from molecules to cells. *Phil Trans R Soc A Math Phys Eng Sci*, 365:2921–2934, 2007.
- [7] Sherwood P, Brooks BR, and Sansom MSP. Multiscale methods for macromolecular simulations. *Curr Opin Struct Biol*, 18:630–640, 2008.
- [8] Jacobs DJ, Rader AJ, Kuhn LA, and Thorpe MF. Protein flexibility predictions using graph theory. *Prot: Struct. Func. Gen.*, 44:150–165, 2001.
- [9] Wells SA, Jimenez-Roldan JE, and Römer RA. Comparative analysis of rigidity across protein families. *Phys. Biol.*, 6(4):046005–11, 2009.
- [10] Wells SA, Menor S, Hespenheide BM, and Thorpe MF. Constrained geometric simulation of diffusive motion in proteins. *Phys. Biol.*, 2:S127–S136, 2005.
- [11] Jolley CC, Wells SA, Hespenheide BM, Thorpe MF, and Fromme P. Docking of photosystem I subunit c using a constrained geometric simulation. *J. Am. Chem. Soc.*, 128(27):8803–8812, 2006.
- [12] Jolley CC, Wells SA, Fromme P, and Thorpe MF. Fitting low-resolution cryo-EM maps of proteins using constrained geometric simulations. *Biophys. J.*, 94:1613–1621, 2008.
- [13] Suhre K and Sanejouand Y-H. ElNémo: a normal mode web server for protein movement analysis and the generation of templates for molecular replacement. *Nucl. Acids Res. (Web Issue)*, 32:610–614, 2004.
- [14] Suhre K and Sanejouand Y-H. On the potential of normal mode analysis for solving difficult molecular replacement problems. *Acta Cryst D*, 60:796–799, 2004.

- [15] Bahar I, Lezon TR, Bakan A, and Shrivastava IH. Normal mode analysis of biomolecular structures: functional mechanisms of membrane proteins. *Chem Rev*, 110:1463–1497, 2010.
- [16] Ma J. Usefulness and limitations of normal mode analysis in modelling dynamics of biomolecular complexes. *Structure*, 13:373–380, 2005.
- [17] Rueda M, Chacón P, and Orozco M. Thorough validation of protein normal mode analysis: A comparative study with essential dynamics. *Structure*, 15:565–575, 2007.
- [18] Nakasako M, Maeno A, Kurimoto E, Harada T, Yamaguchi Y, Oka T, Takayama Y, Iwata, and Kato K. Redox-dependent domain rearrangement of protein disulfide isomerase from a thermophilic fungus. *Biochemistry*, 49:6953–6962, 2010.
- [19] Dykeman EC and Sankey OF. Normal mode analysis and applications in biological physics. *J. Phys. Condens. Matter*, 22:423202, 2010.
- [20] Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, and Bourne PE. The protein data bank. *Nucl. Acids Res.*, 28:235–242, 2000. <http://www.rcsb.org>.
- [21] Wlodawer A, Walter J, Huber R, and Sjolín L. Structure of bovine pancreatic trypsin inhibitor: Results of joint neutron and X-ray refinement of crystal form II. *J. Mol. Biol.*, 180(2):301–329, 1984.
- [22] Amir D and Haas E. Reduced bovine pancreatic trypsin inhibitor has a compact structure. *Biochemistry*, 27(25):8889–8893, 1988.
- [23] K Shipley, M Hekmat-Nejad, J Turner, C Moores, R Anderson, R Milligan, R Sakowicz, and R Fletterick. Structure of a kinesin microtubule depolymerization machine. *EMBO Journal*, 23(7):1422–1432, 2004.
- [24] Elliott PR, Pei XY, Dafforn TR, and Lomas DA. Topography of a 2.0 Å structure of α -1-antitrypsin reveals targets for rational drug design to prevent conformational disease. *Prot. Sci.*, 9(7):1274–1281, 2000.
- [25] Freedman RB, Klappa P, and Ruddock LW. Protein disulfide isomerases exploit synergy between catalytic and specific binding domains. *EMBO Journal*, 15:136–140, 2002.
- [26] Tian G, Kober F, Lewandrowski U, Sickmann A, Lennarz WJ, and Schindelin H. The catalytic activity of protein-disulfide isomerase requires a conformationally flexible molecule. *J. Biol. Phys.*, 283:33630–33640, 2008.
- [27] Tian G, Xiang S, Noiva R, Lennarz WJ, and Schindelin H. The crystal structure of yeast protein disulfide isomerase suggests cooperativity between its active sites. *Cell*, 124:61–73, 2006.
- [28] Hilf RJC and Dutzler R. X-ray structure of a prokaryotic pentameric ligand-gated ion channel. *Nature*, 452(7185):375–379, 2008.
- [29] Word JM, Lovell SC, Richardson JS, and Richardsonzhd DC. Asparagine and glutamine: Using hydrogen atoms contacts in the choice of side-chain amide orientation. *J. Mol. Biol.*, 285:1735–1747, 1999.
- [30] DeLano WL. The PyMOL molecular graphics system. <http://www.pymol.org>, 2002.
- [31] Tirion MM. Large amplitude elastic motions in proteins from single-parameter atomic analysis. *Phys. Rev. Lett.*, 77:1905–1908, 1996.
- [32] Zheng W and Doniach S. A comparative study of motor-protein motions using a simple elastic network model. *Proc. Nat. Acad. Sci.*, 100:13253–58, 2003.
- [33] Cheng X, Lu B, Grant B, Law RJ, and McCammon JA. Channel opening motion of α 7 nicotinic acetylcholine receptor as suggested by normal mode analysis. *J. Mol. Biol.*, 355:310–324, 2006.
- [34] Xy C, Tobi D, and Bahar I. Allosteric changes in protein structure computed by a simple mechanical model: hemoglobin T—R2 transition. *J. Mol. Biol.*, 333:153–168, 2003.
- [35] O Miyashita, JN Onuchic, and PG Wolynes. Nonlinear elasticity, proteinquakes, and the energy landscapes of functional transitions in proteins. *Proc. Nat. Acad. Sci.*, 100:12570–12575, 2003.
- [36] Farrell DW, Speranskiy K, and Thorpe MF. Generating stereochemically-acceptable protein pathways. *Proteins*, 78:2908–2921, 2010.
- [37] Jimenez-Roldan JE, Wells SA, Freedman RB, and RoemerRA. Integration of FIRST, FRODA and NMM in a coarse grained method to study protein disulphide isomerase conformational

- change. In *Journal of Physics: Conference Series*, volume 286, page 012002. IOP Publishing, 2011.
- [38] Krebs WG, Alexandrov V, Wilson CA, Echols E, Yu H, and Gerstein M. Normal mode analysis of macromolecular motions in a database framework: developing mode concentration as a useful classifying statistic. *Proteins*, 48:682–695, 2002.
- [39] Jimenez-Roldan JE, Bhattacharya M, Vishveshwara S, Freedman RB, and Roemer RA. in preperation.
- [40] Belfield W, Cole D, and Payne M. in preparation.
- [41] Burkoff NS, Varnai C, Wells SA, and Wild DL. Exploring the energy landscapes of protein folding simulations with Bayesian computation. *Proc. Nat. Acad. Sci.*, 2011. submitted.

Table 1. The proteins and specific structures selected for this study. For each protein, various E_{cut} values (kcal/mol) were chosen on the basis of rigidity analysis (Figure S1); the Table presents those values used in the simulations of motion that are presented in the main text. Bold values for E_{cut} have been used to compute the xRMSD. Visualization of structures for additional cutoffs (see Table S2) are shown in the supplementary material, see Figures S2 – S7.

Protein	PDB	Resolution	Residues	E_{cut} (kcal/mol)
BPTI	1BPI	1.1Å	58	−0.2, − 2.2
Cytochrome-c	1HRC	1.9Å	105	−0.7, − 1.2
Kinesin	1RY6	1.6Å	360	−0.4, − 1.1
α 1-antitrypsin	1QLP	2.0Å	394	−0.1, − 1.1
PDI	2B5E	2.4Å	504	−0.015, − 0.522 , −1.412
pLGIC	2VL0	3.3Å	1605	−0.4, − 0.5

Figure 1. Schematic of the method. The input (at left) is an all-atom protein structure. Normal mode analysis (above) models the protein with a one-site-per-residue coarse graining and a simple spring model to produce an eigenvector for low-frequency motion. Rigidity analysis (below) identifies noncovalent interactions in an all-atom model of the protein and divides the protein into rigid clusters and flexible linkers. Geometric simulation (right) integrates normal-mode and rigidity information to explore the flexible motion of the protein.

Figure 2. Tertiary structure of all six protein structures (a) BPTI (1BPI), (b) cytochrome-c (1HRC), (c) α 1-antitrypsin (1QLP), (d) kinesin (1RY6,) (e) yeast PDI (2B5E) and (f) pLGIC (2VL0). The structures are given in standard PYMOL [30] format but broken into rigid clusters according to the rigidity analysis (cp. Fig. S1) at the specific values of E_{cut} shown in Table 1. Each rigid cluster is represented in a different colour with the largest rigid cluster indicated in blue and flexible regions shown in black.

Figure 3. Superimposed structural variations and fitted RMSD for small loop motion as found in BPTI and cytochrome-c. Panels (a) and (d) indicate the range of projected tertiary structure for motion along mode m_7 at $E_{\text{cut}} = -2.2$ kcal/mol for BPTI and at $E_{\text{cut}} = -1.2$ kcal/mol for cytochrome-c, respectively. Panels (b,c) and (e,f) show the fitted RMSD as a function of FRODA conformations for BPTI (1BPI) and cytochrome-c (1HRC), respectively, for the non-trivial modes $m_7, \dots, m_{11}$ at two values of E_{cut} as shown. Positive conformation values indicate motion along the direction of the corresponding ELNEMO mode, whereas negative conformation values indicate motion in the opposite direction. Points and error bars indicate mean and standard deviation obtained from five runs of the conformation generation for each mode.

Figure 4. Superimposed structural variation and fitted RMSD for large loop motion as in (a) kinesin (1RY6) and (d) antitrypsin (1QLP) for $E_{\text{cut}} = -1.1$ kcal/mol. Panels (b,c) and (e,f) represent — as in Figure 3 — the fitted RMSD at two values of E_{cut} for kinesin and antitrypsin, respectively. Points and error bars have been determined as in Figure 3.

Figure 5. Superimposed structural variation of large domain motion and fitted RMSD for yeast PDI (2B5E). (a) We show the initial tertiary structure as opaque and the projected structures as partially transparent. All structures are aligned on the central two domains b–b’ to highlight the motion of the a and a’ domains. Motion represents large conformational change along m_7 . Panels (b), (c) and (d) show the fitted RMSDs relative to the initial conformation for three values of E_{cut} . Points and error bars as in Figure 3.

Figure 6. Large scale twist motion in a ligand gated ion channel (2VLO). (a) View down the transmembrane channel in its initial state and projected along m_7 in two directions. (b) Side view showing tilting of the helices during the motion. In both images the structures have been aligned on the extracellular β -sheet portion so as to highlight the relative motion of the domains, and residues from number 283 upwards in each chain are not shown to make the major helices visible. (c and d) Fitted RMSDs relative to initial conformation for low-frequency modes $m_7, \dots, m_{11}$ and two cutoff energies. Points and error bars as in Figure 3.

Figure 7. Extensive RMSD as a function of FRODA conformations for all six proteins moving along mode m_7 . The maximum xRMSD values range from 150Å for BPTI to 5243Å for PDI. There are three clear categories of protein motion: large conformational changes achieved by domain motion (PDI and pLGIC), large loop motions (antitrypsin and kinesin) and small loop motions (BPTI and cytochrome-c). Theselected E_{cut} for each protein are $E_{\text{cut}}^{1\text{BPI}} = -2.2\text{kcal/mol}$, $E_{\text{cut}}^{1\text{HRC}} = -1.2\text{ kcal/mol}$, $E_{\text{cut}}^{1\text{RY6}} = -1.1\text{ kcal/mol}$, $E_{\text{cut}}^{1\text{QLP}} = -1.1\text{ kcal/mol}$, $E_{\text{cut}}^{2\text{B5E}} = -0.522\text{ kcal/mol}$ and $E_{\text{cut}}^{2\text{VL0}} = -0.5\text{ kcal/mol}$. The XRMSD values obtained for $m_8, \dots, m_{11}$ for the selected proteins are consistent with m_7 xRMSD.

Figure 8. Dot product motifs and illustration of mode mixing. (a) Schematic representation of the typical behaviours of the dot product of an initial mode with a current mode during projection along the initial mode eigenvector. Motif 1: gradual, nearly quadratic reduction in the dot product due to a progressive rotation of the current mode compared to the initial one. Eventual constant behaviour indicates that the motions has reached its amplitude limit. Motif 2: more rapid roughly quadratic reduction. Motif 3: sudden collapse of the dot product and the initial mode no longer resembles the current mode with the same mode number. (b) Dot products computed for initial normal mode $m_{11}^{(i)}$ of cytochrome-c and current normal modes from $m_7^{(c)}$ to $m_{26}^{(c)}$. Columns in the positive overlap direction denote projection in the positive direction. Columns in the negative overlap direction indicate projection into the opposite direction. The vertical dashed line marks mode $m_{11}^{(c)}$ and the horizontal dotted line indicates a 0 dot product value.

Supplementary information

S1. Settings, input files and command line options

All simulations were carried out using the ELNEMO and FIRST software. FRODA is a routine available within FIRST. The file formats and command line options to apply a mode eigenvector as a bias in FIRST/FRODA have not previously been published in the literature, and we therefore describe them here.

The required input for all calculations is a `.pdb` format file containing an all-atom representation of the protein structure including hydrogen atoms, which we shall name `protein.pdb`. Our usual procedure is to obtain a file containing the heavy atom positions from the Protein Data Bank; to remove alternate side chain conformations and nonbonded heteroatoms including water molecules; to add hydrogens using the REDUCE software, including flipping of side chains where necessary; and to renumber the atoms sequentially using PYMOL.

Normal mode calculations were carried out using ELNEMO using the default setting in the `pdbmat.dat` input file of a 12Å cutoff in the spring network. The protein structure is given as a `pdbmat.structure` file, consisting of only the α carbon lines from the all-atom structure. The output is a `pdbmat.eigenfacs` file which, for a protein of N residues, contains $3N$ mode eigenvectors. Each eigenvector is described with a mode number, a frequency, and N lines each giving a Cartesian vector; the i th line is the displacement to be applied to the i th residue. The vector is normalized so that the sum of the squares of all displacement vectors is unity. The first few lines of a mode appear thus:

```
VECTOR      7      VALUE  6.0869E-04
-----
 2.5267E-02  2.1069E-02  0.1020
 1.7347E-02  1.5141E-02  9.3303E-02
 3.5897E-02  2.5485E-02  9.2557E-02
...
```

To pass this mode as a bias to FIRST/FRODA we prepare a `mode.in` file giving, for each residue, the identity of the α carbon atom for that residue in the `pdb` file, followed by the displacement vector. The first few lines of a `mode.in` file appear thus:

```
2 0.025267 0.021069 0.102
6 0.017347 0.015141 0.093303
20 0.035897 0.025485 0.092557
...
```

In FRODA, the first displacement vector will be applied as a bias to the motion of all atoms in the first residue, and so forth.

The command line options for FIRST to carry out a FRODA simulation with a mode bias can be given as follows: `FIRST -non $PROTEIN -E -$CUT -FRODA -mobRC1 -freq $FREQ -totconf $TOTCONF -modei -step $STEP -dstep $DSTEP`

Taking terms in order, `FIRST` is the `FIRST` executable; `-non` is to run in noninteractive mode; `$PROTEIN` is the name of the all-atom `.pdb` format input file; `-E -$CUT` runs the rigidity analysis with hydrogen-bond cutoff energy `-$CUT`; `FRODA` invokes the `FRODA` algorithm to simulate flexible motion; `-mobRC1` specifies that the largest rigid cluster is to be mobile, like the smaller rigid clusters, during the simulation; `-freq $FREQ` specifies the frequency with which newly generated conformations are to be written to file as new conformations; `-totconf $TOTCONF` specifies the total number of new conformations to generate; `modei` specifies that `mode.in` should be read and applied as a bias; `-step $STEP` specifies the magnitude of random perturbations of the atomic positions in the generation of each new conformation; and `-dstep $DSTEP` specifies the magnitude of the displacement of the structure along the mode eigenvector in the generation of each new conformation.

For all our calculations we have set `$FREQ` to 100 so as to save every 100th conformation; we have used a `$STEP` of 0.1Å; and we have used `$DSTEP` values of 0.01Å and $-0.01\text{\AA}$ so as to project each mode in both possible directions. Thus to explore the motion of `protein.pdb` at an energy cutoff of -1.0 kcal/mol, we would use two command lines as follows

```
FIRST -non protein.pdb -E -1.000 -FRODA -mobRC1 -freq 100 -totconf 2500
      -modei -step 0.1 -dstep 0.01
```

```
FIRST -non protein.pdb -E -1.000 -FRODA -mobRC1 -freq 100 -totconf 2500
      -modei -step 0.1 -dstep -0.01
```

S2. Raw and fitted RMSD

All the RMSDs reported in the main text are α carbon RMSDs from the input structure to a generated conformation, obtained by least-squares fitting using the `PYMOL intra_fit` command. These values differ somewhat from the raw RMSD values reported by `FRODA` in its output files, which are calculated without any fitting being carried out. In particular, the fitted RMSD saturates once further motion along the mode direction is no longer possible, due to steric clashes or limits imposed by covalent or noncovalent bonding constraints. The raw RMSD is greater than the fitted RMSD and tends not to saturate, but rather to continue to increase slowly, once the motion is effectively jammed. The reason for this different behaviour is a small difference in the statistical weighting given to each residue by `ELNEMO` and by `FIRST/FRODA`. In the elastic network modelling, every residue is given equal statistical weight. The non-trivial mode eigenvectors thus generated have no significant component of rigid-body motion for the whole structure. In `FRODA`, however, the bias is applied to an all-atom representation of the structure; and thus the bias applied to a residue with many atoms affects the whole-body motion of the structure more than the bias applied to a residue with few atoms. The motion in `FRODA` therefore acquires a small component of rigid-body translation and rotation, which increases the raw RMSD. Least-squares fitting to the input structure

*Rapid simulation of protein motion: merging flexibility, rigidity and normal mode analyses*³

removes the rigid-body components, so the fitted RMSD detects actual conformational change.

Supplementary table S1. Extensive RMSD values, maximum RMSD values and the cutoff energies chosen to calculate xRMSD for each protein based on the RCD graphs. For proteins that are expected to be rigid we have chosen a higher E_{cut} and for proteins with an expected conformational change we have chosen a more restrictive E_{cut} .

Protein	Residues	E_{cut} (kcal/mol)	RMSD pos Å	RMSD neg Å	xRMSD pos Å	xRMSD neg Å
BPTI	58	-2.2	2.62	2.66	152	154
Cytochrome-c	105	-1.2	1.44	1.40	151	146
Kinesin	360	-1.1	2.48	2.04	892	733
Antitrypsin	394	-1.1	1.91	2.04	753	804
PDI	504	-0.522	10.54	10.40	5314	5243
pLGIC	1605	-0.5	2.56	2.55	4113	4099

Supplementary table S2. Rigidity analysis cutoff energies. Summary of *all* cutoff energies extracted from the rigidity analysis used in the geometric simulations. RMSD data for only a subset of these are shown in the main text.

Protein	PDB code	Cutoffs (kcal/mol)
BPTI	1BPI	-0.200, -1.700, -2.200
Cytochrome-c	1HRC	-0.700, -1.200
Kinesin	1RY6	-0.400, -0.600, -1.100
Antitrypsin	1QLP	-0.100, -0.500, -1.100
PDI	2B5E	-0.015, -0.522, -0.885, -1.412
pLGIC	2VL0	-0.400, -0.500

Supplementary figure S1. Rigid cluster decomposition graphs for: (a) BPTI (1BPI) (b) cytochrome-c (1HRC) (c) α 1-antitrypsin (1QLP) (d) internal kinesin motor domain (1RY6) (e) yeast PDI (2B5E) and (f) pLGIC (2VL0). The x axis represents the protein backbone and the y axis the energy, E_{cut} , of the last hydrogen bond, which after being removed provokes a change in the rigidity distribution. Each line represents the new rigidity distribution of the polipeptide chain induced by removing a bond which alters the previous rigidity configuration. The residues belonging to rigid clusters are coloured — with the biggest rigid cluster coloured in red, whereas the flexible regions are shown as thin black lines. We choose the energy cutoffs defining the number of rigidity constraints using the RCD plots. For a detailed description of RCD graphs see Refs. [8,9]

Supplementary figure S2. Tertiary structure of BPTI (1BPI). Colouring is defined using the rigidity analysis results shown in Figure S1. Flexible regions are illustrated in black whereas rigid residues are coloured as per the rigid cluster they belong to. The biggest rigid cluster is coloured in blue. The number and size of the rigid clusters vary depending on the chosen cutoff value, which for BPTI are (a) $E_{\text{cut}} = -0.2$ kcal/mol, (b) $E_{\text{cut}} = -1.7$ kcal/mol and (c) $E_{\text{cut}} = -2.2$ kcal/mol. Note that the colour code used to represent residues within the same rigid cluster is not the same in the RCD and in the tertiary structures. The biggest rigid cluster in the RCD graphs is noted in red and in the tertiary structures is noted in blue.

Supplementary figure S3. Tertiary structure of cytochrome-c (1HRC). Colouring of the tertiary structure is defined as in Figure S2 but with cutoff energies (a) $E_{\text{cut}} = -0.7$ kcal/mol and (b) $E_{\text{cut}} = -1.2$ kcal/mol.

Supplementary figure S4. Tertiary structure of α 1-antitrypsin (1QLP). Colouring of the tertiary structure is defined as per Figure S2 but at cutoff energies of (a) $E_{\text{cut}} = -0.1$ kcal/mol, (b) $E_{\text{cut}} = -0.5$ kcal/mol and (c) $E_{\text{cut}} = -1.1$ kcal/mol.

Supplementary figure S5. Tertiary structure of internal kinesin motor domain (1RY6). Colouring of the tertiary structure is defined as in Figure S2 but with cutoff energies (a) $E_{\text{cut}} = -0.4$ kcal/mol, (b) $E_{\text{cut}} = -0.6$ kcal/mol and (c) $E_{\text{cut}} = -1.1$ kcal/mol.

Supplementary figure S6. Tertiary structure of yeast PDI (2B5E). Colouring of the tertiary structure is defined as in Figure S2 but with cutoff energies (a) $E_{\text{cut}} = -0.015$ kcal/mol, (b) $E_{\text{cut}} = -0.522$ kcal/mol, (c) $E_{\text{cut}} = -0.885$ kcal/mol and (d) $E_{\text{cut}} = -1.412$ kcal/mol. Note that colors are assigned according to cluster size which changes depending on the cutoff energy.

Supplementary figure S7. Tertiary structure of pLGIC (2VL0). Colouring of the tertiary structure is defined as in Figure S2 but with cutoff energies (a) $E_{\text{cut}} = -0.4$ kcal/mol and (b) $E_{\text{cut}} = -0.5$ kcal/mol. Note that the protein appears to be rigid for $E_{\text{cut}} = -0.4$ kcal/mol and that there is a switch like rigidity (second order) transition at $E_{\text{cut}} = -0.5$ kcal/mol which reveals the most flexible parts of the secondary structure which allow mobility.

Supplementary figure S8. Dot product graph for BPTI (1BPI). The dot product $m_j^{(i)} \cdot m_j^{(c)}$ between an initial starting mode $m_j^{(i)}$ and its current mode $m_j^{(c)}$, $j = 7, \dots, 11$ as the initial structure is projected along the initial mode. The current modes, $m_j^{(c)}$, are obtained from performing normal mode analysis on the current conformations as the initial structure is projected along an initial mode $m_j^{(i)}$. For clarity, dot products for only 25 conformations of each direction of motion are shown. The evolution of the dot product along the conformations is reported for different cutoff energies, which for BPTI are (a) $E_{\text{cut}} = -0.2$ kcal/mol, (b) $E_{\text{cut}} = -1.7$ kcal/mol and (c) $E_{\text{cut}} = -2.2$ kcal/mol. The horizontal and vertical dotted lines denote the largest possible value of $m_j^{(i)} \cdot m_j^{(c)}$ and the zero on the conformer axis, respectively.

Supplementary figure S9. Dot product graph for cytochrome-c (1HRC) as described in Figure S8 but with E_{cut} values of (a) -0.7 kcal/mol and (b) -1.2 kcal/mol.

Supplementary figure S10. Dot product graph for $\alpha 1$ -antitrypsin (1QLP) as described in Figure S8 but with E_{cut} values of (a) $E_{\text{cut}} = -0.1$ kcal/mol, (b) $E_{\text{cut}} = -0.5$ kcal/mol and (c) $E_{\text{cut}} = -1.1$ kcal/mol.

Supplementary figure S11. Dot product graph for internal kinesin motor domain (1RY6) as described in Figure S8 but with E_{cut} values of (a) -0.4 kcal/mol (b) -0.6 kcal/mol and (c) -1.1 kcal/mol.

Supplementary figure S12. Dot product graph for yeast PDI (2B5E) as described in Figure S8 but with E_{cut} values of (a) $E_{\text{cut}} = -0.015$ kcal/mol, (b) $E_{\text{cut}} = -0.522$ kcal/mol, (c) $E_{\text{cut}} = -0.885$ kcal/mol and (d) $E_{\text{cut}} = -1.412$ kcal/mol.

Supplementary figure S13. Dot product graph for a ligand gated ion channel protein (2VL0) as described in Figure S8 but with E_{cut} values of (a) $E_{\text{cut}} = -0.4$ kcal/mol and (b) at $E_{\text{cut}} = -0.5$ kcal/mol.