

Validating and improving elastic network models with molecular dynamics simulations

Tod D. Romo and Alan Grossfield*

Department of Biochemistry and Biophysics, University of Rochester Medical Center, Rochester, New York 14642

ABSTRACT

Elastic network models (ENMs) are a class of simple models intended to represent the collective motions of proteins. In contrast to all-atom molecular dynamics simulations, the low computational investment required to use an ENM makes them ideal for speculative hypothesis-testing situations. Historically, ENMs have been validated via comparison to crystallographic B-factors, but this comparison is relatively low-resolution and only tests the predictions of relative flexibility. In this work, we systematically validate and optimize a number of ENM-type models by quantitatively comparing their predictions to microsecond-scale all-atom simulations of three different G protein coupled receptors. We show that, despite their apparent simplicity, well-optimized ENMs perform remarkably well, reproducing the protein fluctuations with an accuracy comparable to what one would expect from all-atom simulations run for several hundred nanoseconds.

Proteins 2011; 79:23–34.
© 2010 Wiley-Liss, Inc.

Key words: normal modes; principal component analysis; convergence; fluctuations; G protein-coupled receptors.

INTRODUCTION

Molecular modeling is an important method for understanding the structure and dynamics, and hence function, of biological systems. This is particularly important for systems such as membrane proteins where isolation and biophysical characterization are difficult. The obvious choice for modeling the dynamics of a system is all-atom molecular dynamics (MD), which represents the motions and interactions of the system with atomic detail. While this method can provide extraordinarily detailed information, it is computationally very expensive. Microsecond-scale simulations of G protein-coupled receptors (GPCRs) embedded in a lipid bilayer represent the state of the art in large-system long time-scale simulations and require supercomputing resources or specialized hardware.^{1–6}

Approaches that capture the predictive power of MD, at least on an abstract level, and yet are tractable using readily-available computational power would be valuable. Elastic network models (ENMs) are one promising approach. Network models are based on the observation that the low-frequency modes, corresponding to large-scale collective motions are relatively insensitive to the details of the underlying model in normal mode analysis (NMA). Early work by Tirion⁷ found that replacing the detailed forcefield typically used in a NMA with single-parameter harmonic potentials yielded modes virtually indistinguishable from the more detailed models. ENMs expand upon this idea by replacing the detailed forcefield used in NMA and MD with a set of elastic springs connecting a network of nodes that represent the topology of the protein system. A thorough review of the history and application of ENMs can be found in Bahar *et al.*⁸ In particular, Bahar's group has applied ENMs to GPCRs with significant success.^{9,10}

Despite the apparent simplicity of the models used in ENMs, their ability to describe collective motions is impressive, particularly when compared with their computational cost. The typical network model consists of a set of nodes attached to each other by simple harmonic springs with minimum energy distances set to match the starting structure. Most commonly, C α 's are used as the nodes; all C α 's within a prescribed distance (12–15 Å) are connected to each other. Higher level abstractions for nodes are also possible and are often found in multiscale modeling methods.^{8,11} Conversely, there is no constraint on using more detailed models, such as C α nodes with additional nodes to represent side-chains; the objective is to balance the predictive power and computational expense.

The dimensionality of the problem can also be reduced by partitioning the model, and hence the Hessian.^{12,13} One such method, called Vibrational Subsystem Analysis^{14,15} (VSA), partitions the model into a region of interest (subsystem) for which a NMA is computed, and an environment that can react to the subsystem modes.

*Correspondence to: Alan Grossfield, Department of Biochemistry and Biophysics, University of Rochester Medical Center, Rochester, NY, 14642. E-mail: alan_grossfield@urmc.rochester.edu.

Received 16 April 2010; Revised 13 July 2010; Accepted 31 July 2010

Published online 19 August 2010 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/prot.22855




Figure 1

The VSA network for β_2 AR. The $C\alpha$ nodes are shown as gray spheres. The lines connecting the nodes represent the ENM “springs.” The springs that are colored red are part of the “environment” while the blue springs are the “subsystem” composed of the transmembrane helices. For illustrative purposes, a distance cutoff of 10 Å was used in determining spring connections.

In the context of comparing ENM results for different proteins, the VSA partitioning provides a convenient method for reducing the NMA to the common subset between proteins while retaining the contributions from the extraneous parts. In this work, we compare three different GPCR structures, each composed of a seven-transmembrane helical bundle with intervening intra- and extracellular loops of different lengths. We partition the protein into a subsystem composed of the common transmembrane (TM) $C\alpha$'s and relegate the remaining $C\alpha$'s (including the loops) to the environment. Such a partitioning is depicted in Figure 1 for the β_2 -adrenergic receptor (β_2 AR), using a cutoff radius of 10 Å to improve the clarity of the figure. The alternative approach is to clip the network to the common nodes. The potential difficulty here is that modes at the edges of the subset (e.g., the ends of the TM helices) may be perturbed since their interactions with the rest of the protein are absent.

In this work, we examine the use of long time-scale molecular dynamics simulations for validating and improving

upon the ENM results. We use VSA to compare the collective motions of different GPCRs and contrast this with the more “classical” anisotropic network model (ANM). We also show that improving the spring constants used in the ENM can have a significant impact on the ability of the ENM to reproduce the fluctuations found in the long time-scale molecular dynamics. Finally, we give a method for assessing the time-scale equivalence between the ENM and MD fluctuations, that is, the length of an MD simulation required to predict the fluctuations in a long simulation with equal fidelity as the ENM.

METHODS

Three different microsecond scale all-atom molecular dynamics simulations were used in this study. The first is a 1.02 μ s simulation of the β_2 AR embedded in a lipid bilayer via the constant temperature ensemble.¹ The second system is a 1.6 μ s simulation of dark-state rhodopsin embedded in a bilayer using the constant energy ensemble.⁴ The final system is a ≈ 1.9 μ s simulation of an activated (protonated) Cannabinoid CB2 receptor (CB2) embedded in a bilayer using the NVT ensemble.¹⁶ All simulations were performed at an average temperature of 310 K using BlueMatter⁶ on the BlueGene/W supercomputer⁵ located at the IBM T. J. Watson Research Center in Yorktown Heights, NY. In the subsequent analysis, all MD trajectories were sampled at 1 ns.

Elastic network models

ENM construction

The ENM was constructed by taking the $C\alpha$'s to be the nodes of the network. For each protein, we computed the average $C\alpha$ position from the entire trajectory using an iterative alignment procedure¹⁷ with the TM $C\alpha$ atoms as reference points. An illustration of what the network looks like for β_2 AR is shown in Figure 1.

The Hessian used is the typical ANM Hessian^{18–20}:

$$H_{ij} = \gamma \Gamma_{ij} \begin{bmatrix} \frac{\partial^2 R_{ij}}{\partial X_i \partial X_j} & \frac{\partial^2 R_{ij}}{\partial X_i \partial Y_j} & \frac{\partial^2 R_{ij}}{\partial X_i \partial Z_j} \\ \frac{\partial^2 R_{ij}}{\partial Y_i \partial X_j} & \frac{\partial^2 R_{ij}}{\partial Y_i \partial Y_j} & \frac{\partial^2 R_{ij}}{\partial Y_i \partial Z_j} \\ \frac{\partial^2 R_{ij}}{\partial Z_i \partial X_j} & \frac{\partial^2 R_{ij}}{\partial Z_i \partial Y_j} & \frac{\partial^2 R_{ij}}{\partial Z_i \partial Z_j} \end{bmatrix}, i \neq j \quad (1)$$

$$H_{ii} = - \sum_{j, j \neq i} H_{i,j}$$

where H_{ij} is a 3×3 superblock of the Hessian, Γ_{ij} is the spring constant between the i th and j th nodes, and γ is a constant used to assign an effective temperature. R_{ij} is the vector representing the spring, that is, between the two nodes' equilibrium position.

In the ANM formulation, the spring constants are uniform, that is, $\Gamma_{ij} = 1$ for $R_{ij} < r_c$ and 0 otherwise. The

HCA method²¹ instead uses distance-based spring constants:

$$\Gamma(R_{ij}) = \begin{cases} ar + b & \text{for } r < r_c \\ cr^{-d} & \text{for } r \geq r_c \end{cases} \quad (2)$$

where r is the length of the intranode vector R_{ij} . The parameters originally used in HCA were $r_c = 4.0 \text{ \AA}$, $a = 205.5 \text{ kcal mol}^{-1} \text{ \AA}^{-3}$, $b = 571.2 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$, $c = 3.059 \times 10^5 \text{ kcal mol}^{-1} \text{ \AA}^4$, and $d = 6$. In our implementation, the overall cutoff radius used in constructing the network was 15 \AA . In addition, the spring constants were constrained to always be positive (i.e., negative spring constants, found at very short distances, were set to 0).

Vibrational subsystem analysis

The idea behind VSA is to partition the system into a subsystem and an environment, integrating out the “environment” fluctuations to report their average effect on the subsystem.^{12,14,15} The environmental degrees of freedom are integrated out while allowing the environment to respond to changes in the subsystem by minimizing the total energy. Although a full derivation is available in Woodcock *et al.*,¹⁴ we will sketch the underlying logic here. The normal mode equation is rewritten with the partitioning as,

$$\begin{bmatrix} H_{ss} & H_{se} \\ H_{es} & H_{ee} \end{bmatrix} \begin{bmatrix} V_s \\ V_e \end{bmatrix} = \lambda \begin{bmatrix} M_s & 0 \\ 0 & M_e \end{bmatrix} \begin{bmatrix} V_s \\ V_e \end{bmatrix} \quad (3)$$

H_{ss} and H_{ee} are the subsystem and environment Hessians respectively, and M_s and M_e are diagonal matrices containing the masses for the corresponding nodes, in triplicate to match the dimensionality of the Hessians.

The normal mode equations above can be reduced to a generalized eigenvalue problem,

$$H'_{ss} V_s = \lambda M'_s V_s \quad (4)$$

where the “effective” subset Hessian and mass matrices, H'_{ss} and M'_s are given by,

$$\begin{aligned} H'_{ss} &= H_{ss} - H_{se} H_{ee}^{-1} H_{es} \\ M'_s &= M_s + H_{se} H_{ee}^{-1} M_e H_{ee}^{-1} H_{es} \end{aligned}$$

The resulting subsystem eigenvectors, V_s , are “mass orthogonal” and must be mass-weighted to be orthogonal in cartesian space and comparable between different solutions:

$$V'_s = M_s^{1/2} V_s \quad (5)$$

We used the LAPACK²² routine DSYGVX to solve the generalized eigenproblem above, and DTRMM to perform the Cholesky decomposition²³ of M_s and hence the square root.

The contribution of the environment’s mass can be excluded by setting M_e to zero, leaving only the diagonal matrix M_s in Eq. (4) above.^{12,24} In our implementation, the subsystem used unit masses and the corresponding eigenvectors, V_s , are orthogonal and do not require mass-weighting. We used DGESVD to compute the singular value decomposition (SVD) of H'_{ss} .

In this study, we investigated three different methods of handling mass in VSA. We used the zero environment mass with unit masses for the subsystems. We also assigned unit masses to both the subsystem and the environment. Finally, we assigned the total residue mass to each C α node.

To apply the VSA method to GPCR’s, each GPCR structure was partitioned into a subsystem, consisting of the common TM C α ’s, and an environment, consisting of the non-TM C α ’s (i.e., the loop regions). To simplify comparison, we selected the equivalent residues from each protein based on a sequence alignment: 1.30–1.59 (in Ballesteros-Weinstein notation²⁵), 2.38–2.67, 3.22–3.54, 4.40–4.62, 5.35–5.68, 6.29–6.60, and 7.33–7.55. This yielded a subsystem composed of 205 C α ’s, out of 282 total for β_2 AR, 360 for CB2, and 384 for rhodopsin.

We also re-optimized the five constants in the HCA method— a , b , c , d , and r_c —using a Nelder-Mead simplex optimizer²⁶ to maximize the covariance overlap between the ENM and all-atom molecular dynamics principal component analysis (PCA) (see Section “Principal component analysis” for details), subject to the constraints that r_c and d must always be positive and nonzero. In the case of fitting against multiple models, that is, simultaneously optimizing both β_2 AR and rhodopsin, the sum of the covariance overlaps (see Section “Covariance and subspace overlap”) for each model was maximized. CB2 was left out of the optimization as a control. Each variant of VSA was optimized for a new set of HCA parameters, referred to as HCA*, and the corresponding parameters are shown in Table I.

Implementation

The different ENM models and spring constant assignments are implemented using the LOOS library²⁷ and are freely available from SourceForge (<http://loos.sourceforge.net>). LOOS is an object-oriented library implemented in C++ that uses Boost and ATLAS (for linear algebra) to provide a lightweight but powerful library designed for facilitating the creation of new analytical tools for molecular dynamics. Included in LOOS is a “selection expression” parser that enables tools to easily select atoms out of a model for analysis based on user-provided selections. LOOS supports the native file formats for most major MD packages, such as CHARMM²⁸/NAMD,²⁹ Amber,³⁰ Gromacs,³¹ and Tinker.³² LOOS comes bundled with over 50 different tools, including those for computing different ENMs.

Table 1
HCA Constants Optimized Against β_2 AR and Rhodopsin via Simplex for Different VSA Implementations

Method	r_c (Å)	a (kcal mol ⁻¹ Å ⁻³)	b (kcal mol ⁻¹ Å ⁻²)	c (kcal mol ⁻¹ Å ⁴)	d
$M_e = 0$	3.479	235.246	821.952	5.340×10^5	10.025
$M_s = M_e = 1$	4.148	275.042	908.567	5.332×10^5	7.146
Residue mass	4.304	283.160	933.608	5.366×10^5	7.100

$M_e = 0$ used unit subsystem and zero environment masses. $M_s = M_e = 1$ used unit masses for both partitions. “Residue mass” assigned the total mass of the residue to the ENM node.

Mode comparison

Principal component analysis

PCA of a molecular dynamics trajectory yields a set of eigenvectors that point along the modes of motion for each atom. We implemented PCA in LOOS²⁷ using the SVD,²³ an alternative methodology that is equivalent to the more common covariance matrix analysis formalism. Briefly, at each time point, the chosen coordinates are stacked vertically to form a $3n$ dimensional vector (where n is the number of atoms) representing the protein conformation at that time point. These column “conformation” vectors are concatenated to form a $3n \times L$ matrix A , where L is the length of the simulation. Thus, A represents the ensemble of structures in the simulation. The average conformation is subtracted off, $\hat{A} = A - \bar{A}$, where the average is computed by an iterative alignment scheme.¹⁷ This eliminates rigid body rotations and translations in the MD simulation. Finally, the SVD is computed,

$$A = U\Sigma V^T \quad (6)$$

using LAPACK.²² The U matrix is a unitary orthonormal matrix whose columns, called left singular vectors, describe the modes of motion of atoms in A and are equivalent to the eigenvectors of $A^T A$ in the more common PCA formalism. The diagonal elements of Σ , or σ_b are the singular values of A and are equivalent to $\sqrt{\lambda_i}$, where λ_i are the eigenvalues of the cross-correlation matrix $A^T A$.

Covariance and subspace overlap

We used two different measures to compare the subspaces determined by both the ENM and PCA methods. The first is the subspace overlap,^{33,34} defined as

$$\Psi_{A,B} = \frac{1}{N} \sum_i^N \sum_j^N (\vec{v}_i^A \cdot \vec{v}_j^B)^2 \quad (7)$$

where \vec{v}_i^A is the i th eigenvector from simulation A and \vec{v}_j^B is the j th eigenvector from simulation B . It is typically computed for a subset of eigenvectors, that is, $N < 3n$ where N is the number of modes, since it is usually the first few modes that contribute most to the overall fluctua-

tions of the system. The value ranges from 0, where the modes are completely dissimilar to 1, where the modes are identical (or more precisely, the subspace spanned by each set of modes are identical). Critically, the eigenvalues are not used.

By contrast, the covariance overlap measures not just the similarity of the eigenvectors, but their relative importance as well.^{17,33,34} The covariance overlap ranges in value from 0 (no overlap between the fluctuation space) to 1 (identical fluctuation spaces), and is defined as,

$$\Omega_{A,B} = 1 - \left[\frac{\sum_i^{N_{\text{modes}}} (\lambda_i^A + \lambda_i^B) - 2 \sum_i^{N_{\text{modes}}} \sum_j^{N_{\text{modes}}} \sqrt{\lambda_i^A \lambda_j^B} (\vec{v}_i^A \cdot \vec{v}_j^B)^2}{\sum_i^{N_{\text{modes}}} (\lambda_i^A + \lambda_i^B)} \right]^{1/2} \quad (8)$$

The eigenvalues obtained from an ENM represent frequency, whereas the PCA eigenvalues represent amplitude. In the harmonic model, $\omega \propto 1/\lambda$, therefore to compare the ENM eigenvalues with the PCA eigenvalues, we must use the reciprocal ENM eigenvalues. Both approaches give six zero eigenvalues however, representing the rigid-body rotation and translation of the entire system, which must be excluded. Next, because ENMs only set relative magnitudes of motion, they must be scaled such that the total power (or magnitude of fluctuations) between the ENM and PCA results are similar. This is performed by scaling the inverted ENM eigenvalues by,

$$k = \frac{\sum_i \lambda_i^A}{\sum_i \lambda_i^B} \quad (9)$$

where λ_i^A is the i th PCA eigenvalue and λ_i^B is the i th reciprocal ENM eigenvalue.

Block overlap

To assess the ability of a shorter simulation to predict the motions described by a longer trajectory, a “block overlap” method was devised, similar in spirit to block averaging.^{35,36} Given a long trajectory, a PCA is first computed for the entire length, L frames, of the simulation; for the purposes of this analysis, the full trajectory is treated as the “gold standard.” The trajectory is then divided into N contiguous blocks of L/N frames. A PCA

is computed for each block, and the covariance overlap [Eq. (8)] with the full trajectory is computed. The average overlap is then plotted as a function of increasing block size. As the blocks get longer, the covariance overlap will asymptotically approach 1.

RESULTS AND DISCUSSION

Several groups have explored different ways to set the individual spring constants including assigning a different constant for backbone atoms,³⁷ to different constants for covalent versus noncovalent interactions,³⁸ and a “parameter free” method that does not use a cutoff and uses an inverse square distance as a spring constant.³⁹ Other groups have investigated using short (tens of nanoseconds) MD simulations to derive improved spring constants.^{40,41}

The particular method we tested is known as the HCA method,²¹ shown in Eq. (2), which uses two different distance-based functions selected by a cutoff radius. The standard ANM Hessian for β_2 AR and the HCA Hessian are shown in Figure 2. It is readily apparent that in spite of using a 15 Å cutoff for the ANM Hessian, there is a sharp cutoff between neighboring nodes and more distant nodes. In contrast, the HCA Hessian shows both a narrower range in constants and a more uniform distribution of values. Panel C shows the Hessian using an optimized set of HCA constants (HCA*, described above) that sharpens the Hessian and increases the dynamic range.

Because we used the MD average structure as the model for the ENM, which may be “unphysical,” it is possible that nodes may be sufficiently close together that they are assigned negative spring constants by the HCA-based methods. This can lead to negative eigenvalues and a failure of the underlying NMA. We therefore constrained the spring constants such that any negative constants were set to 0. An alternative strategy is to find the structure in the trajectory that is closest to the MD average and to use this for the network model. We implemented this and found that the overall trends described here held, but the covariance overlaps are slightly lower than those found with the MD average and truncated spring constants.

Method comparison

Comparing the subspaces defined by the different models and methods requires comparing not only the directions defined by the eigenvectors but also the shape of the fluctuation space defined by the eigenvalues as well. One common measure is the subspace overlap shown in Eq. (7). The subspace overlap can be driven arbitrarily to 1 by choosing a sufficiently large cutoff n , so the cutoff is usually chosen such that the corresponding modes describe the bulk of the fluctuations in the model (i.e., the first several dozen modes). The difficulty

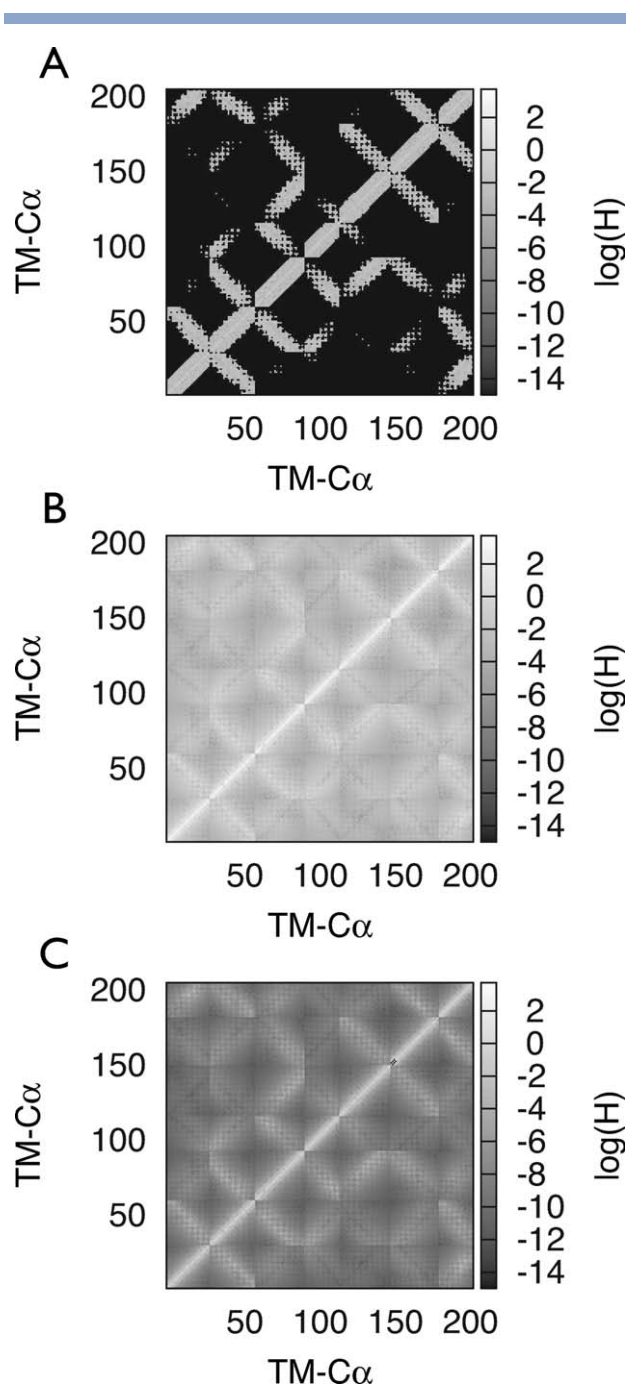


Figure 2

Effect of different methods of assigning spring constants on the β_2 AR ANM Hessian matrix. Panel A shows the standard ANM method using a cutoff radius of 15 Å. Panel B shows the HCA method. Panel C shows the simplex-optimized HCA constants. This illustrates the increase in dynamic range of the Hessian in the simplex-optimized HCA method and how both methods include a broader range of neighbor-contacts over the standard ANM Hessian.

is that it now depends upon an arbitrary cutoff, and it still does not consider the shape of the corresponding power spectrum—if mode 1 contributes 12% of the motion in one solution, but the corresponding mode in

Table IICovariance and Subspace ($N = 25$) Overlap Between Different Network Models and Principal Components Computed from Molecular Dynamics

Methods	β_2AR		Rhodopsin		CB2	
	Covariance	Subspace	Covariance	Subspace	Covariance	Subspace
ANM	0.37	0.37	0.32	0.33	0.36	0.40
ANM (HCA)	0.50	0.54	0.42	0.46	0.46	0.50
ANM (HCA*)	0.55	0.66	0.44	0.51	0.47	0.57
VSA [†]	0.36	0.43	0.32	0.45	0.35	0.41
VSA [†] (HCA)	0.48	0.60	0.41	0.54	0.43	0.56
VSA [†] (HCA*)	0.56	0.69	0.48	0.55	0.52	0.61
VSA [‡]	0.37	0.42	0.33	0.43	0.34	0.40
VSA [‡] (HCA)	0.47	0.55	0.41	0.52	0.42	0.50
VSA [‡] (HCA*)	0.49	0.56	0.42	0.49	0.42	0.49
VSA [§]	0.37	0.42	0.32	0.41	0.34	0.38
VSA [§] (HCA)	0.47	0.53	0.40	0.50	0.42	0.49
VSA [§] (HCA*)	0.48	0.54	0.42	0.49	0.41	0.47

*Optimized HCA parameters. The bold overlaps indicate CB2 was not used to optimize the HCA parameters.

[†]VSA computed using zero masses for the environment.[‡]VSA computed using unit masses.[§]VSA computed using sum of all masses in the corresponding residue.

another solution is the fifth and only contributes 3%, this difference in their relative importance is lost. Alternatively, imagine a two-dimensional fluctuation space where the principal axes are identical, but where one distribution is a narrow “pencil” shape and the other is nearly circular. The eigenvectors for these spaces match each other, giving a subspace overlap of 1. The shapes described by the eigenvectors and eigenvalues, however, are vastly different. In contrast, the covariance overlap [Eq. (8)] weights the dot products by the corresponding eigenvalues incorporating information about the shape of the respective power spectra. An important consideration when using the covariance overlap occurs when both sets of the eigenvalues are all similar in magnitude. In this case, the covariance overlap turns into effectively a subspace overlap over the entire set of eigenvectors, leading to an unexpectedly high overlap. The covariance overlap also reduces to the subspace overlap in the case where all λ_i are either 0 or 1.

The comparison of the different methods for the three different systems against their long time-scale MD simulation PCA results is shown in Table II. Both the covariance and subspace overlaps are shown. For each method, the “standard” Hessian along with the HCA and HCA* Hessians are compared. The benefit of a distance-based spring constant assignment are obvious. The use of HCA in either ANM or VSA outperforms the standard method. Applying a simple optimization to the HCA constants yields a significant improvement in the massless VSA overlaps. Interestingly, optimizing the HCA parameters for β_2AR and rhodopsin also improves the overlap for the massless VSA of CB2 even though CB2 was not used in the optimization. The overlaps for CB2 do not change when using the other forms of VSA (each ENM method was optimized individually). An additional

cross-validation of the optimizations was performed where the HCA parameters were optimized for β_2AR then the VSA HCA* was computed for rhodopsin and compared against the MD PCA result. The resulting covariance overlap was 0.46. Similarly, optimizing against rhodopsin and using these HCA parameters for β_2AR resulted in a covariance overlap of 0.54. The resulting HCA parameters were slightly different between the two individual optimizations, but the general form (i.e., $3.0 \text{ \AA} \leq r_c \leq 3.5 \text{ \AA}$ and $d = 10$) was the same. This suggests that the optimization against long time-scale MD may convey a universal improvement in the ENMs, at least across class A GPCRs. This result strongly suggests that the predictive power of ENMs in general can be improved through careful choice of their spring constants.

Surprisingly, the VSA implementations with mass did not outperform the ANM models. Indeed, even the best massless VSA model performed only slightly better than the ANM with HCA*. It is possible that this is due to the particular systems explored here; perhaps, the edge effects are not as large for simple helical bundles like GPCRs. Regardless, VSA also suggests a clean formalism for including membrane-protein interactions in an ENM approach, which may be important for further research.

One of the more interesting features of the β_2AR trajectory is the partial unwinding of the cytoplasmic end of TM helix 6 (TM6) with the opening of the ionic lock and the subsequent reforming as the lock closes,¹ which is captured by the first mode of the PCA on that trajectory, shown in Figure 3. Searching the massless VSA HCA* modes for the closest match in the top-25 modes finds the 20th mode with a dot-product of 0.82. In contrast, the best match in the top-25 modes for the massless VSA HCA is the 25th mode with a dot-product of 0.5 (not shown in Fig. 3).




Figure 3

The first mode of the PCA for β_2 AR is shown here as dark gray vectors focusing on the intracellular end of TM6. In the MD simulation, this stretch of TM6 unwinds briefly during the opening of the ionic lock, and then reforms as the lock closes. The 20th mode of the massless VSA HCA* method is shown in light gray. No similar vectors were found in the top 25-modes of the nonoptimized VSA HCA results. This shows a qualitative improvement in the biological relevance of the modes predicted by the VSA HCA* method.

Comparing GPCRs

The PCA of the MD trajectories show that the fluctuations between the different GPCRs are not very similar. Table III lists the similarity of fluctuation spaces between the different GPCRs for massless VSA, massless VSA with HCA*, and PCA on the entire trajectory. The PCA results were compared by normalizing the eigenvalues by the trajectory length. The ENM results were compared without using any scaling [i.e., Eq. (9)] to test whether the ENM captured differences in overall flexibility. The low covariance and subspace overlaps for the PCA between the different GPCRs however indicates that the fluctuations sampled by each are not very similar.

Surprisingly, VSA overstates the covariance overlap to a remarkable degree, although VSA-HCA and VSA-HCA* do not. Because the subspace overlap does not show the same problem, we believe this phenomenon is a due to peculiarities of the eigenvalue distribution for VSA, in particular the significant contributions from the long “tail” of relatively similar, nonzero eigenvalues. We tested this numerically by creating an artificial data set consisting of the eigenvectors generated from PCA of β_2 AR combined with eigenvalues from VSA. The covariance overlap of this synthetic data set with VSA eigenvalues and eigenvectors was 0.74. By contrast, if we reverse the

comparison, using a synthetic data set consisting of VSA eigenvectors and PCA eigenvalues with the true PCA data set, gives a far more reasonable overlap of 0.36. This suggests that it is the VSA eigenvalues that are driving the anomalous covariance overlap.

A further numerical test demonstrates that the tail plays the critical role. We fit the PCA eigenvalue distribution to a double exponential, $f(x) = a \exp(-x/b) + c \exp(-x/kd) + e$ and generated multiple sets of eigenvalues where $1 \leq k \leq 10$. The covariance overlap was then computed using the eigenvectors from the VSA and PCA models of β_2 AR with the simulated eigenvalues. The covariance overlap ranged from 0.39 for $k = 1$ to 0.65 for $k = 10$. In other words, by increasing the relative contribution of the tail of the eigenvalue distribution, we can drive the covariance overlap higher.

Since the subspace overlap only considers the first N modes and neglects the eigenvalues, it is not vulnerable to this difficulty, at the cost of containing less information. Were the fluctuations more similar between the different GPCRs with VSA, we would have expected the subspace overlap to be higher for VSA than the more refined models. Instead, the better ENM models (e.g., VSA-HCA*), show increased similarity with the results from PCA for both overlap quantities. For example, the covariance overlap for VSA HCA between β_2 AR and rhodopsin is 0.57 with a subspace overlap of 0.31 ($N = 25$). The performance of VSA HCA* (shown in Table III) is even closer.

In addition, the distribution of fluctuation magnitudes along the PCA modes are quite different between these GPCRs. The average mean squared fluctuation along each mode is shown in Figure 4. This is defined as $\frac{\lambda_i}{nL}$ where λ_i is the i th eigenvalue, n is the number of atoms, and L is the trajectory length (i.e., number of structures).

Figure 4 shows that β_2 AR is significantly less mobile (smaller fluctuations) than rhodopsin and, for the first mode, is half that of CB2. The rigidity of β_2 AR relative to rhodopsin was previously noted in Romo *et al.*¹ We see here that the increased motion is confined to the 10 lowest-frequency modes for rhodopsin, and 15 or so modes for CB2. The magnitude of fluctuations for the higher-frequency modes are comparable amongst the GPCRs.

Table III

Comparison of Different ENM Methods and PCA Between the Different GPCRs

	Rhodopsin			CB2		
	VSA	VSA HCA*	PCA	VSA	VSA HCA*	PCA
β_2 AR	0.70/0.23	0.41/0.36	0.26/0.33	0.70/0.25	0.41/0.34	0.26/0.32
Rhodopsin				0.71/0.28	0.42/0.34	0.24/0.29

The first number is the covariance overlap followed by the subspace overlap ($N = 25$). The VSA used zero masses for the environment subset.

*Optimized HCA parameters.

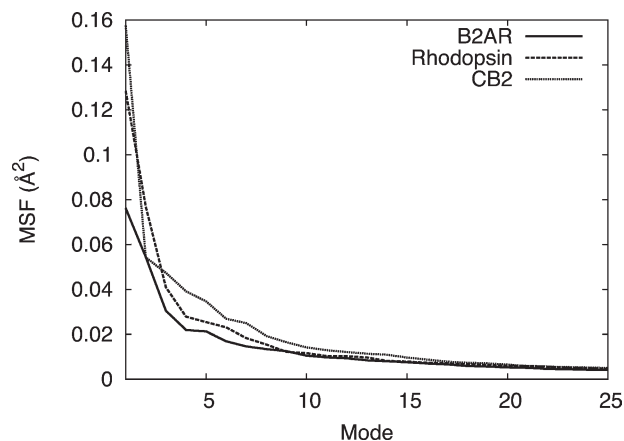


Figure 4

The average mean-square fluctuations per atom contributed by each PCA mode is shown for each GPCR. β_2 AR is clearly more rigid than rhodopsin and CB2. Beyond 10 modes, the difference is negligible.

Power spectra

The results from the covariance overlap emphasize that comparing the directions of the eigenvectors is necessary but not sufficient to fully compare the two fluctuation spaces. The shape of the distribution of modes, that is, their power spectra must also be taken into consideration. The power spectra for both ANM, massless VSA, and PCA for rhodopsin are shown in Figure 5. The fractional contribution of each mode to the total fluctuations, that is, $\frac{\lambda_i}{\sum_i \lambda_i}$, is plotted on a log-scale. Here again the PCA of the MD trajectory is taken as the “gold standard.” Figure 5(A) shows the spectrum over all 609 modes. It is immediately obvious that the ANM and VSA models underestimate the magnitude of the low-frequency modes while severely overestimating the magnitude of the higher-frequency modes. Although the VSA-HCA model improves the spectrum slightly, it is still considerably different from the PCA spectrum. The ANM-HCA* and VSA-HCA* methods, in contrast, are much closer to the PCA spectrum, although they still underestimate the importance of the first mode. Interestingly, the PCA spectrum shows a small “dip” near the mode 400, and all HCA methods show a sigmoidal step at nearly the same mode.

Figure 5(B) shows a closeup of the first 30 modes. Here, the underestimation of the magnitude of the first 15–30 modes for the ANM, VSA, and VSA-HCA methods are obvious.

This under-over estimation pattern seen with the ENM models can also be found when comparing shorter MD simulations against a longer one. The time-dependent PCA power spectrum, shown for rhodopsin in Figure 6, is computed by dividing the MD trajectory into contiguous blocks of a given size, performing PCA on each block, and computing the corresponding fractional con-

tribution to fluctuations for each mode. These are then averaged over all blocks of the same size and plotted against block-size. Block sizes up to 250 ns, already a substantial length for a system of this size, underestimate the contribution of the first two modes by around 35% and then overestimate the higher-frequency modes. Even at 750 ns, the lower frequency modes are underestimated and the higher frequency modes overestimated. This figure shows that care must be taken when making assertions about collective motions based on short MD simulations and that very long trajectories are needed if MD is used to guide ENM enhancements.^{17,35} Similarly, it is possible that the microsecond-scale simulations used here are long enough to characterize the low-frequency collective fluctuations, while very long by current standards.

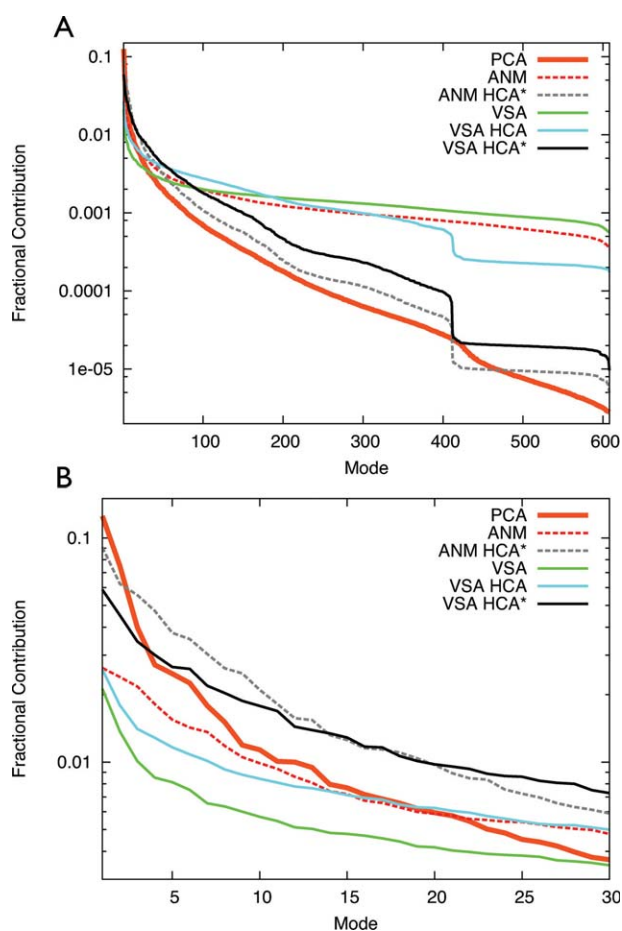


Figure 5

Panel A shows the rhodopsin power spectrum between PCA, VSA, and ANM using the standard of assigning spring constants, the HCA method, and the optimized HCA method. The power spectrum is normalized to 1, giving the contribution of each mode to the overall motion of the model, and plotted on a log scale. Panel B shows a closeup of the first 30 modes. A consistent pattern of underestimating the magnitude of the low-frequency motions while overestimating the high frequency motions is seen in all ENMs. However, a significant improvement is seen in the HCA* power spectra.




Figure 6

The PCA power spectrum depends on the length of the trajectory. The first 10 modes of the average power spectrum for different block sizes is shown. The power spectrum for the full trajectory PCA is shown in black. As the trajectory gets longer, more importance is assigned to the lowest frequency modes. Mode 2 also changes significantly from 750 to 1605 ns.

ENM robustness

The ability of the ENM solution to withstand perturbations in the equilibrium structure used for the network is of some concern. It is possible that by using the average structure from the MD simulation, we are biasing our ENM results more favorably toward the MD PCA. We assessed the ability of the ENM to withstand perturbations by the following method: first, the average structure for the entire trajectory is found using an optimal iterative alignment procedure.¹⁷ This structure is used to compute a reference ENM using the massless VSA HCA*. Frames are drawn from the MD trajectory at random and used to compute a new ENM, and the covariance overlap is then computed between the two ENM results. In Figure 7, the covariance overlaps are plotted against the RMSD between the frame and the average structure. In addition, the covariance overlap for ENMs constructed from five different X-ray crystallographic structures of rhodopsin are also shown. This shows that there is considerable overlap between ENM solutions where the starting structure is within 1.5 Å RMSD. Even at the 2 Å RMSD of the crystal structures, there is still a significant overlap. The similarity in fluctuation spaces between the crystal structure ENMs and the average MD structure ENM does translate into similar overlaps with the MD PCA. Computing a massless VSA-HCA* for each crystal structure gives a covariance overlap of 0.46 for all of the crystal structures, compared with 0.48 for the MD average structure.

Comparing efficiency of ENM and MD

All-atom molecular dynamics is the gold standard for modeling protein fluctuations. However, it is computa-

tionally very expensive and sampling limitations lead to statistical errors. By contrast, ENMs use a simple model but have no statistical error. The question becomes, how long must an MD simulation run before the statistical error is smaller than the systematic error of the ENMs? The first step is to quantify the statistical error present in the MD. We do this by constructing a “block overlap” as described in Section “Block overlap.” The average covariance overlap as a function of block size for β_2 AR is shown in Figure 8(A) and for rhodopsin in Figure 8(B). Block sizes range from 25 ns to half of the trajectory (500 and 800 ns, respectively). In the cases where the block size is smaller than the number of degrees of freedom, there will be multiple zero eigenvalues. The covariance overlap ignores these, by construction, and the number of modes used in the subspace overlap is always smaller than the number of nonzero eigenvalues. Treating the full-length trajectory as the “gold standard,” this graph then shows how well a shorter simulation can be expected to produce “correct” fluctuations. Convergence occurs very slowly; the β_2 AR covariance overlap does not reach 0.5 until the blocks are 200 ns long, and even at 500 ns, the overlap is only 0.7. Rhodopsin requires nearly 500 ns to reach an overlap of 0.5 and only approaches 0.6 at the 800 ns point. This is not surprising considering how much more flexible rhodopsin is than β_2 AR (see Fig. 4). Moreover, while these trajectories are long, by current standards, they almost certainly still have significant statistical error.³⁵

The covariance overlap for the PCA of a given block size, compared against the full MD trajectory, gives an estimate for what time-scale the ENM is equivalent to in




Figure 7

The sensitivity of ENM solutions to structural perturbations is shown by comparing massless VSA HCA* ENMs computed from randomly selected frames of the full rhodopsin trajectory against the ENM computed using the average structure from the entire trajectory. The covariance overlap between five different X-ray crystal structures and the overall average ENM is also shown. This demonstrates that the ENM solutions are relatively insensitive to perturbations in the starting structure.




Figure 8

The covariance overlap for β_2 AR (panel A), rhodopsin (panel B), and CB2 (panel C). The block overlap is the average overlap for all blocks of a given size with the modes from the full trajectory. The error bars are the standard deviations of the overlap for all blocks of a given size. The maximum block size is half the length of the trajectory. The horizontal lines illustrate the overlap values for the original VSA model, giving an estimate for what time-scale the ENM is equivalent to in terms of describing the fluctuations in a “gold standard” (long time-scale) MD simulation. The VSA model using HCA spring constants and the VSA model using the simplex optimized HCA constants. All VSA models used unit subsystem masses and zero environment masses. CB2 was not used in determining the HCA* parameters.

terms of the ENM’s ability to describe the fluctuations present in a long time-scale MD simulation. For example, the massless VSA overlap for β_2 AR is 0.36, which intersects the block overlap curve at 50 ns. Thus, we can say that the sampling quality of VSA is roughly equivalent to 50 ns of simulation; considering the relative computational expense of the two methods, this is impressive. Similarly, the VSA-HCA overlap is 0.48, corresponding to a block size of about 150 ns. Improving the spring constants in the Hessian has increased the effective “simulation time” of the ENM by 100 ns. The VSA-HCA* is 0.56, which corresponds to another 100 ns increase. The best ENM model for β_2 AR then is equivalent to roughly 250 ns of MD simulation, yet computing this result takes seconds on a contemporary desktop as opposed to weeks on a supercomputer. Interestingly, in the case of rhodopsin, the effective simulation times are even longer with VSA equivalent to 100 ns, VSA-HCA equivalent to 225 ns, and VSA-HCA* equivalent to 400 ns of simulation time. This occurs because rhodopsin is more flexible than β_2 AR and thus sampling it with MD is more challenging.

Because CB2 was not used in optimizing the HCA parameters used in the VSA calculations for Figure 8, it serves as a control.

Figure 8(C) shows the fluctuations predicted by the VSA-HCA* of CB2 to be equivalent to ~ 600 ns of MD. This is double the equivalent time-scale for the standard HCA implementation. This further suggests the optimizations made in HCA* using β_2 AR and rhodopsin are not specific to either those two systems or their MD trajectories.

Long time-scale molecular dynamics simulations are invaluable for understanding the low-frequency collective motions of biomolecules. This comes at a substantial cost however. A quarter of a microsecond of all-atom MD for the systems studied here takes several weeks to run on a supercomputer. In contrast, the massless VSA HCA* takes less than 5 s on a modern desktop, yet its ability to accurately reproduce the low-frequency collective motions of a multi-microsecond simulation is no worse than the month-long MD. This is even more remarkable considering the simplicity of the model (e.g., ignoring side-chains) and that ENMs assume a harmonic model. One caveat is that there is far more information available in the MD simulation. However, it is arguable that if one is interested primarily in larger scale motions, then one might as well use ENM.

CONCLUSIONS

We have compared several different ENM methods with long time-scale all-atom MD simulations for three different GPCR systems. We found that while VSA provides a convenient means for separating similar structures into comparable and different partitions, it does not per-

form significantly better than more traditional ENMs (i.e., ANM) with respect to predicting the fluctuations found in the MD simulations.

Our block overlap analysis and time-dependent power spectrum analysis show that all-atom simulations of a length that would normally be considered “long,” on the order of several hundred nanoseconds, do not necessarily describe the low-frequency fluctuations of the longer simulations accurately. Indeed, there is a systematic error present in both the shorter MD simulations and the ENMs, where the low-frequency motions are underestimated and the higher-frequency motions overestimated.

We have shown that by optimizing the spring constants used in constructing the ENM Hessian, we can significantly improve the ENM’s ability to reproduce the fluctuations found in long MD simulations. The effect on the ENM is equivalent to simulating an additional 100–300 ns of a comparable MD simulation. Moreover, these optimizations are not specific to β_2 AR and rhodopsin, as shown by the improvements in CB2, which was not part of the optimization procedure. Examining the power spectra, however, indicates that while we have improved the model, there is likely still room for improvement with more sophisticated spring constants models.

Finally, we have shown that the ENMs are fairly robust with respect to the structure used to build the network. These findings have important implications on the future development and use of ENMs for probing long time-scale collective motions. They should also serve as a reminder that care must be taken when characterizing the low frequency collective motions with MD.

ACKNOWLEDGMENTS

We thank Nick Leioatts for pioneering the VSA implementation in LOOS. We thank David Mathews, Harry Stern, Nick Leioatts, and Joshua Horn for their discussions and critical reading of this manuscript. We also thank An Ghysels and Wenjun Zheng for helpful discussions regarding VSA. We also thank the IBM Watson Research Center and the Computational Biology Center for use of the BGW BlueGene supercomputer. We extend a special thanks to the Blue Matter development team (B. Fitch, R. Germain, A. Rayshubskiy, T. J. C. Ward, M. Eleftheriou, F. Suits, Y. Zhestkov, R. Zhou, J. Pitera, and W. Swope).

REFERENCES

- Romo TD, Grossfield A, Pitman MC. Concerted interconversion between ionic lock substates of the β -2 adrenergic receptor revealed by microsecond timescale molecular dynamics. *Biophys J* 2010; 98:76–84.
- Dror RO, Arlow DH, Borhani DW, Jensen MØ, Piana S, Shaw DE. Identification of two distinct inactive conformations of the beta2-adrenergic receptor reconciles structural and biochemical observations. *Proc Natl Acad Sci USA* 2009;106:4689–4694.
- Shaw D, Deneroff M, Dror R, Kuskin J, Larson R, Salmon J, Young C, Batson B, Bowers K, Chao J, Eastwood M, Gagliardo J, Grossman J, Ho C, Ierardi D, Kolossváry I, Klepeis J, Layman T, McLeavey C, Moraes M, Mueller R, Priest E, Shan Y, Spengler J, Theobald M, Towles B, Wang S.Anton, a special-purpose machine for molecular dynamics simulation. ISCA '07: proceedings of the 34th annual international symposium on computer architecture, San Diego, CA. 2007. pp. 1–12.
- Grossfield A, Pitman MC, Feller SE, Soubias O, Gawrisch K. Internal hydration increases during activation of the G-protein-coupled receptor rhodopsin. *J Mol Biol* 2008;381:478–486.
- Allen F, Almasi G, Andreoni W, Beece D, Berne BJ, Bright A, Brunheroto J, Cascaval C, Castanos J, Coteus P, Crumley P, Curioni A, Denneau M, Donath W, Eleftheriou M, Fitch B, Fleisher B, Georgiou CJ, Germain R, Giampapa M, Gresh D, Gupta M, Haring R, Ho H, Hochschild P, Hummel S, Jonas T, Lieber D, Martyna G, Maturu K, Moreira J, News D, Newton M, Philhower R, Picunko T, Pitera J, Pitman M, Rand R, Royyuru A, Salapura V, Sanomiya A, Shah R, Sham Y, Singh S, Snir M, Suits F, Swetz R, Swope WC, Vishnumurthy N, Ward TJC, Warren H, Zhou R. Blue gene: a vision for protein science using a petaflop supercomputer. *IBM Syst J* 2001;40:310.
- Fitch BG, Germain RS, Mendell M, Pitera J, Pitman M, Rayshubskiy A, Sham Y, Suits F, Swope WC, Ward TJC, Zhestkov Y, Zhou R. Blue Matter, an application framework for molecular simulation on blue gene. *J Parallel Distrib Comput* 2003;63:759–773.
- Tirion M. Large amplitude elastic motions in proteins from a single-parameter, atomic analysis. *Phys Rev Lett* 1996;77:1905–1908.
- Bahar I, Lezon TR, Bakan A, Shrivastava IH. Normal mode analysis of biomolecular structures: functional mechanisms of membrane proteins. *Chem Rev* 2010;110:1463–1497.
- Rader AJ, Anderson G, Isin B, Khorana HG, Bahar I, Klein-Seetharaman J. Identification of core amino acids stabilizing rhodopsin. *Proc Natl Acad Sci USA* 2004;101:7246–7251.
- Isin B, Rader AJ, Dhiman HK, Klein-Seetharaman J, Bahar I. Pre-disposition of the dark state of rhodopsin to functional changes in structure. *Proteins* 2006;65:970–983.
- Jang H, Na S, Eom K. Multiscale network model for large protein dynamics. *J Chem Phys* 2009;131:245106.
- Ghysels A, Speybroeck VV, Pauwels E, Catak S, Brooks BR, Neck DV, Waroquier M. Comparative study of various normal mode analysis techniques based on partial Hessians. *J Comput Chem* 2010; 31:994–1007.
- Ghysels A, Speybroeck VV, Pauwels E. Mobile block hessian approach with adjoined blocks: an efficient approach for the calculation of frequencies in macromolecules. *J Chem Theory Comput* 2009;5:1203–1215.
- Woodcock HL, Zheng W, Ghysels A, Shao Y, Kong J, Brooks BR. Vibrational subsystem analysis: a method for probing free energies and correlations in the harmonic limit. *J Chem Phys* 2008;129:214109.
- Hafner J, Zheng W. Approximate normal mode analysis based on vibrational subsystem analysis with high accuracy and efficiency. *J Chem Phys* 2009;130:194111.
- Hurst DP, Grossfield A, Lynch DL, Feller S, Romo TD, Gawrisch K, Pitman MC, Reggio PH. A lipid pathway for ligand binding is necessary for a cannabinoid G protein-coupled receptor. *J Biol Chem* 2010;285:17954–17964.
- Grossfield A, Feller SE, Pitman MC. Convergence of molecular dynamics simulations of membrane proteins. *Proteins* 2007;67:31–40.
- Eyal E, Yang LW, Bahar I. Anisotropic network model: systematic evaluation and a new web interface. *Bioinformatics* 2006;22:2619–2627.
- Doruker P, Atilgan AR, Bahar I. Dynamics of proteins predicted by molecular dynamics simulations and analytical approaches: application to alpha-amylase inhibitor. *Proteins* 2000;40:512–524.
- Atilgan AR, Durell SR, Jernigan RL, Demirel MC, Keskin O, Bahar I. Anisotropy of fluctuation dynamics of proteins with an elastic network model. *Biophys J* 2001;80:505–515.

21. Hinsen K, Petrescu A, Dellerue S. Harmonicity in slow protein dynamics. *Chem Phys* 2000;261:25–37.
22. Anderson E, Bai Z, Dongarra J, Greenbaum A, McKenney A, Du Croz J, Hammerling S, Demmel J, Bischof C, Sorensen D. Lapack: a portable linear algebra library for high-performance computers. In: *Supercomputing '90: proceedings of the 1990 conference on Supercomputing*. Los Alamitos, CA, USA: IEEE Computer Society Press; 1990. pp. 2–11.
23. Stewart GW. *Introduction to matrix computations*. Computer science and applied mathematics. New York, NY: Academic Press, Inc; 1973.
24. Zheng W, Brooks BR. Probing the local dynamics of nucleotide-binding pocket coupled to the global dynamics: myosin versus kinesin. *Biophys J* 2005;89:167–178.
25. Ballesteros JA, Weinstein H. Integrated methods for the construction of three-dimensional models and computational probing of structure-function relations in G protein-coupled receptors. *Methods Neurosci* 1995;25:366–428.
26. Nelder JA, Mead R. A simplex method for function minimization. *Comput J* 1964;7:308–313.
27. Romo TD, Grossfield A. Loos: an extensible platform for the structural analysis of simulations. *Conf Proc IEEE Eng Med Biol Soc* 2009;1:2332–2335.
28. Brooks B, Bruccoleri R, Olafson B, States D, Swaminathan S, Karplus M. Charmm: a program for macromolecular energy, minimization, and dynamics calculations. *J Computat Chem* 1983;4:187–217.
29. Phillips JC, Braun R, Wang W, Gumbart J, Tajkhorshid E, Villa E, Chipot C, Skeel RD, Kalé L, Schulten K. Scalable molecular dynamics with NAMD. *J Comput Chem* 2005;26:1781–802.
30. Case DA, Cheatham TE, Darden T, Gohlke H, Luo R, Merz KM, Onufriev A, Simmerling C, Wang B, Woods RJ. The amber biomolecular simulation programs. *J Comput Chem* 2005;26:1668–1688.
31. Spoel DVD, Lindahl E, Hess B, Groenhof G, Mark AE, Berendsen HJC. Gromacs: fast, flexible, and free. *J Comput Chem* 2005;26:1701–1718.
32. Ponder J. Tinker 4.2. Available at: <http://dasher.wustl.edu/tinker>.
33. Hess B. Convergence of sampling in protein simulations. *Phys Rev E Stat Nonlin Soft Matter Phys* 2002;65:031910.
34. Faraldo-Gómez JD, Forrest LR, Baaden M, Bond PJ, Domene C, Patargias G, Cuthbertson J, Sansom MSP. Conformational sampling and dynamics of membrane proteins from 10-nanosecond computer simulations. *Proteins* 2004;57:783–791.
35. Grossfield A, Zuckerman DM. Quantifying uncertainty and sampling quality in biomolecular simulations. *Annu Rep in Comput Chem* 2009;5:23–48.
36. Flyvbjerg H, Petersen H. Error estimates on averages of correlated data. *J Chem Phys* 1989;91:88–103.
37. Ming D, Wall ME. Allostery in a coarse-grained model of protein dynamics. *Phys Rev Lett* 2005;95:198103.
38. Kondrashov DA, Cui Q, Phillips GN. Optimization and evaluation of a coarse-grained model of protein motion using x-ray crystal data. *Biophys J* 2006;91:2760–2767.
39. Yang LW, Song G, Jernigan R. Protein elastic network models and the ranges of cooperativity. *Proc Natl Acad Sci USA* 2009;106:12347–12352.
40. Moritsugu K, Smith JC. Coarse-grained biomolecular simulation with reach: realistic extension algorithm via covariance hessian. *Biophys J* 2007;93:3460–3469.
41. Lyman E, Pfaendtner J, Voth GA. Systematic multiscale parameterization of heterogeneous elastic network models of proteins. *Biophys J* 2008;95:4183–4192.