



Published in final edited form as:

J Chem Theory Comput. 2012 July 10; 8(7): 2424–2434. doi:10.1021/ct3000316.

Elastic Network Models are Robust to Variations in Formalism

Nicholas Leioatts, Tod D. Romo, and Alan Grossfield*

Department of Biochemistry & Biophysics, University of Rochester Medical Center, Rochester, NY 14642, USA

Abstract

Understanding the functions of biomolecules requires insight not only from structures, but from dynamics as well. Often, the most interesting processes occur on time scales too slow for exploration by conventional molecular dynamics (MD) simulations. For this reason, alternative computational methods such as elastic network models (ENMs) have become increasingly popular. These simple, coarse-grained models represent molecules as beads connected by harmonic springs; the system's motions are solved analytically by normal mode analysis. In the past few years, many different formalisms for performing ENM calculations have emerged, and several have been optimized using all-atom MD simulations. In contrast to other studies, we have compared the various formalisms in a systematic, quantitative way. In this study, we optimize many ENM functional forms using a uniform dataset containing only long ($> 1 \mu\text{s}$) all-atom MD simulations. Our results show that all models once optimized produce spring constants for immediate neighboring residues that are orders of magnitude stiffer than more distal contacts. In addition, the statistical significance of ENM performance varied with model resolution. We also show that fitting long trajectories does not improve ENM performance due to a problem inherent in all network models tested: they underestimate the relative importance of the most concerted motions. Finally, we characterize ENMs' resilience by tessellating the parameter space to show that broad ranges of parameters produce similar quality predictions. Taken together our data reveals that choice of spring function and parameters are not vital to performance of a network model and that simple parameters can be derived "by hand" when no data is available for fitting, thus illustrating the robustness of these models.

1 Introduction

Structural changes in proteins are vital to activation and regulation, and many proteins function via conformational rearrangements that involve displacing large pieces of the molecule. These concerted or global motions often occur on the microsecond to millisecond time scale, making their study by all-atom computational methods intractable to all but a few research groups with highly specialized resources.^{1–3} However, insight into protein dynamics is valuable, and even approximate interactions can help guide experiments and tailor hypotheses. For this reason, elastic network models (ENMs) have become very popular.⁴ ENMs are a class of simple models for protein motion, where the system is represented as a network of beads inter-connected by harmonic springs with lengths matching those found in the protein's native state; usually, the system is coarse-grained by only representing the C_α 's of a protein. By applying normal mode analysis (NMA) to this

*To whom correspondence should be addressed: alan_grossfield@urmc.rochester.edu.

Supporting Information

Supporting information includes definitions of all spring functions used in this study and details of the opsin simulation. In addition, the results of our standard score and covariance overlap analysis are available in tables S1–S3 and the parameters obtained from fitting all three proteins in aggregate are reported in table S4. This information is available free of charge via the internet at <http://pubs.acs.org/>.

model, one can extract the collective motions of the system with minimal computational investment. Network models have garnered much attention due to their success in predicting global motions.^{5–13} They can also be a practical aid in designing and addressing the feasibility of experiments; for example, these models have proven useful in evaluating the functional motions of proteins^{6–9,14–16} as well as refining structural data.^{17–19} Notably, because the normal modes of motion are solved analytically, there is no statistical error in their predictions, a problem that generally plagues molecular dynamics.²⁰

ENMs' ease of implementation and computational efficiency makes them a suitable starting point for many structural studies; however, their utility is limited in some areas due to their simplicity.^{19,21,22} To counter this lack of detail, numerous groups have developed features designed to augment the predictive power of ENMs.^{23–34} We explore several such methods of improving ENMs while utilizing long molecular dynamics (MD) trajectories to evaluate each model's efficacy. Specifically, we examine the performance of different spring functional forms and resolutions, systematically analyzing their performance after fitting to simulation. For this study we will use the anisotropic form of the network model first described by Atilgan et al.²³

Historically, the performance of ENMs has been validated by comparison to crystallographic B-factors. However, this is a limited comparison since B-factors at best describe the flexibility of different regions of the proteins, usually isotropically. By contrast, the ENM-NMA provides far more information, describing the set of coherent motions available to the system, as well as their relative amplitudes. For this reason, we prefer to compare to molecular dynamics simulations, applying principal component analysis (PCA) to the trajectories to extract analogous information;³⁵ the eigenvectors have an analogous meaning in PCA and ENM-NMA. The eigenvalues have different meanings — magnitude in PCA and frequency in ENM-NMA — but given the harmonic nature of ENMs, one can simply invert the frequencies obtained from ENM-NMA to give a unitless measure of the relative magnitude of motion. In the current study, we will compare harmonic NMA and quasi-harmonic PCA eigendecompositions^{15,36,37} (details of this comparison are outlined in Section 2.2).

Recently, there have been many articles aimed at refining the functional form of the spring contacts in network models using MD. This is a worthwhile undertaking as all-atom MD is considered the gold-standard among molecular simulation techniques. However, in each case short simulations were used, leading to concern about the statistical quality of the data^{20,35,38–40} as well as systematic errors.³⁵ For this evaluation, we chose to use a smaller parametrization set of three proteins for which we had long trajectories available (each greater than 1 μ s in length). We hypothesized that fitting long trajectories was vital since it is precisely the slow motions that we extract from the NMA of ENMs. In a recent publication, however, Liu et al. suggested that simulations spanning a range of two orders of magnitude produce similar dominant motions.⁴¹ This is an interesting result that cannot be overlooked; however, they consider only the most dominant motion, as opposed to comparing the whole fluctuation space, which was done in the current study. We previously showed that fitting to short simulations leads to systematic errors which manifest in the eigenvalue spectra.³⁵ This result will be further explored in the present manuscript.

We also evaluate the effects of varying model resolution. ENMs were originally developed as a lower-resolution alternative to all-atom NMA.^{37,42,43} Each residue is typically represented by one bead placed at the C_{α} position, but occasionally system size or level of detail restrains this decision. Several models exist for altering resolution,^{19,26,44} but for this study we decided on two simple variations to extract a trend. While performance varied, the

new resolutions were capable of predicting motions with a similar fidelity to the standard model.

2 Methods

Many variants of the basic ENM method have been published, including a variety of functional forms to define the spring constants,^{24,25,27,30–33,45} several levels of resolution,^{19,26,44} and methods of interfacing with the surrounding environment.²¹ There have also been many attempts to refine these models based on crystallographic B-factors^{16,28,46} and by comparison to all-atom MD.^{30–33,35,36} In the present work we parametrize spring functions using only long MD simulations.

2.1 Elastic Network Models

ENMs are constructed from a network of beads connected by Hookean springs. These beads are typically located at the C_α 's of each residue when studying proteins. In the current study, average structures were calculated from the MD trajectories using an iterative alignment.⁴⁷ Formally, this network of interactions is described by the following potential:

$$U_{ij}=k(R_{ij}^o) (|R_{ij}|-|R_{ij}^o|)^2 \quad (1)$$

where R_{ij} is the distance between two particular beads, i and j , R_{ij}^o is their equilibrium distance, and $k(R_{ij}^o)$ is a variable spring function. The input structure is designated as the relaxed position of all the springs, making it the global minimum by definition. The functional form of $k(R_{ij}^o)$, or the spring function, can be defined by one of several formalisms. In addition to the descriptions below, explicit definitions of all spring functions can be found in the supplemental text.

In the standard ENM, this spring function simply identifies connections. Two beads are designated as connected if they are within a cutoff distance, R_c , and not connected otherwise. Formally, this is defined by the Heaviside function:

$$k(R_{ij}^o)=\begin{cases} 1 & :R_{ij}^o < R_c \\ 0 & :R_{ij}^o \geq R_c \end{cases} \quad (2)$$

This is the simplest spring function, but it has a few significant weaknesses. First, we intuitively expect that the interactions between distant residues should be weaker than interactions between adjacent ones. Although this is captured in an approximate way by the use of a simple cutoff, one could arguably improve the behavior by dispensing with the cutoff and making the spring constant a smoothly varying function of the minimum energy distance, R_{ij}^o . Moreover, this resolves a numerical instability found in the pure cutoff approach: if the cutoff is too short then certain beads may not have enough interactions resulting in underdetermined dynamics. In that case, the resulting eigendecomposition will fail.²³

One example of a spring that varies smoothly is an exponentially distance-dependent function:

$$k(R_{ij}^o)=\exp(-aR_{ij}^o) \quad (3)$$

Here a is an adjustable scaling factor. This function connects all beads in the model while weighting the “stiffness” of the spring by the distance between beads. Several network models use this form (in part), including those introduced by Hinsén,⁴⁸ Kovacs,²⁵ Smith,³⁰ Voth,³¹ Orozco,³³ and Yang.³²

A similar, more detailed model was proposed by Hinsén and coworkers.²⁴ In this formalism, referred to as HCA, beads closer than the distance R_c are connected by a linearly distance-dependent spring, while beads outside the cutoff are scaled by a power decay:

$$k(R_{ij}^o) = \begin{cases} a(R_{ij}^o)+b & :R_{ij}^o < R_c \\ c(R_{ij}^o)^{-d} & :R_{ij}^o \geq R_c \end{cases} \quad (4)$$

Here the parameters a, b, c , and d in addition to R_c are all adjustable.

Other methods explicitly take the chemical structure of the molecule into account by using different spring functions for virtual “bonded” interactions — those between residues adjacent in sequence — than for “non-bonded” interactions. An example similar to that of Hinsén et al. would be:

$$k(R_{ij}^o) = \begin{cases} k_1 & : \text{Bonded} \\ c(R_{ij}^o)^{-d} & : \text{Non-bonded} \end{cases} \quad (5)$$

Here, one other small difference was introduced; instead of scaling “bonded” terms by distance they are all given the same stiffness, k_1 . In addition, multiple levels of connectivity can be implemented. For instance, the REACH³⁰ method uses virtual bonds, angles, and torsions, connecting first, second, and third nearest neighboring C_α 's, respectively, followed by a non-bonded term.

Any of the above spring functions can be applied to Eq. (1), defining the potential energy between two beads. This is done for all pairs of beads in the system and the resulting Hessian is then computed. The normal modes of this system can be solved by an eigendecomposition; the eigenvectors represent the directions of motion, and the eigenvalues represent their associated frequencies. A more detailed description of this diagonalization has been outlined previously (see for example refs. 4,23,49). In this work, all network models were implemented using the LOOS “ElasticNetworks” package (freely available at <http://loos.sourceforge.net>).^{50,51} LOOS is a lightweight structure analysis library that reads all major MD model and trajectory file formats.

2.1.1 Varying the Resolution of ENMs—While network models are generally abstracted to C_α resolution, we implemented and validated a higher and lower resolution as well. These can be used to fit the scope of a particular project, either looking for more subtle details¹⁹ or predicting the motions of a very large system.^{6–8,14}

Our higher resolution model incorporates a second bead per residue positioned at each side chain's center of mass (CoM) as previously employed by Bahar & Jernigan⁵² and Zheng.¹⁹ The lower resolution model uses a single bead for every two residues placed at the center of mass of the pair. Throughout this manuscript we will refer to these as the two-bead per residue model and one-bead per two-residue model, respectively.

We applied each of the spring functions discussed above to all three resolutions. For models that take connectivity into account, we included bonds between neighboring C_α 's and from each C_α to its side chain bead.

2.2 Validating Against Molecular Dynamics Simulations

To evaluate the normal modes of the network models, we compared them to microsecond-scale all-atom MD of three G protein-coupled receptors (GPCRs): rhodopsin (RHOD in figure legends), β_2 adrenergic receptor (β_2 AR), and cannabinoid receptor 2 (CB2). These simulations were published separately and were also used in our prior ENM studies.^{35,53–56} Briefly, the rhodopsin simulation contained 49 1-stearoyl-2-docosahexaenoyl-phosphatidylcholine (SDPC), 50 1-stearoyl-2-docosahexaenoyl-phosphatidylethanolamine (SDPE), and 24 cholesterol molecules and was conducted using the NVE ensemble. The simulation of β_2 AR contained 99 1-palmitoyl-2-oleoyl-phosphatidylcholine (POPC) lipids and was conducted using the NVT ensemble. The CB2 simulation contained 123 POPC and 38 sn-2-arachidonoylglycerol molecules and was conducted using the NVT ensemble. All of these simulations were performed using the CHARMM27r force field for the lipids,^{57,58} and the CHARMM22 force field for the protein.⁵⁹

In addition, simulations of two other proteins were used just for testing our results. Specifically, we analyze a pair of $\approx 1 \mu\text{s}$ trajectories of the A2A adenosine receptor (A2A) in a POPC bilayer, run by Lyman and coworkers.⁶⁰ For this analysis, we concatenated the two simulations, for a total simulation time of $2.2 \mu\text{s}$. Additionally, we considered a 472 ns trajectory of opsin (the apo form of bovine rhodopsin), containing 123 SDPE lipids; further details of the simulation can be found in the supplemental text.

As in previous work, we performed principal component analysis (PCA) on the trajectories taking advantage of the fact that PCA produces information analogous to that acquired from ENMs. While the PCA of a simulation depends on the covariance matrix and is explicitly not harmonic, its similarity to harmonic analyses like NMA have been extensively studied in the literature.^{15,36,37} The eigenvectors have analogous physical meanings in the two formalisms: the direction of motion associated with a particular eigenvalue. The eigenvalues produce similar, although not identical information. In the case of simulation PCA, they represent the relative amplitude of motion as sampled from a thermalized trajectory, while in network model NMA they represent the relative frequencies of the motion. Because ENMs are explicitly harmonic in nature, the frequencies can simply be inverted to get effective amplitudes. In light of the similarities between these two analyses, we will use the quasi-harmonic PCA of a simulation to parametrize ENM-NMA results. We compare the results of PCA and ENM-NMA using the same methods used to compare two PCA results. In particular, we used the covariance overlap, because it uses the eigenvalues to account for the relative contribution of each eigenvector.^{35,38}

$$\Omega_{A,B} = 1 - \left[\frac{\sum_i^{3N} (\lambda_i^A + \lambda_i^B) - \sum_i^{3N} \sum_j^{3N} \sqrt{\lambda_i^A \lambda_j^B} (\vec{v}_i^A \cdot \vec{v}_j^B)^2}{\sum_i^{3N} (\lambda_i^A + \lambda_i^B)} \right]^{1/2} \quad (6)$$

Here λ represents the eigenvalue associated with the eigenvector, \mathbf{v} . For the ENM-NMA normal modes, the eigenvalues were inverted to produce quantities proportional to relative amplitude. The superscripts A and B denote the two results being compared while the subscripts i and j indicate this summation is over all $3N-6$ modes (excluding the three purely

translational and three purely rotational modes which have zero frequency). The inverted eigenvalues are normalized across a single eigendecomposition prior to application of Eq. (6), such that the total amplitude of predicted motion is identical. This metric scales from 0 to 1, where 0 indicates that the eigenvectors are orthogonal and 1 indicates complete overlap of the eigenspace. It is worth noting that this measure reduces to something close to the subspace overlap⁶¹ if one assumes the eigenvalue spectrum contains only ones and zeroes. The parameters obtained from our fitting (using a Nelder-Mead Simplex⁶² optimization) are documented (for all resolutions and spring functions) in Supplementary Table S4.

2.2.1 Estimation of Equivalent MD Trajectory Length—The physical significance of a particular covariance overlap value is difficult to interpret in the absence of context. Therefore, in order to better understand the comparison between ENM-NMA results and PCA, we compare the performance of our ENM-NMA results to the sampling of an MD trajectory. To do this, we combined the covariance overlap with a block averaging approach as in previous work.^{35,40} Block averaging allows assessment of the sampling quality of a small contiguous piece of the trajectory relative to the full simulation. Briefly, this entails splitting the trajectory into many contiguous blocks and computing the principal components of each block individually. The average covariance overlap between the principal components of these blocked trajectories and the full simulation PCA describes how well the short simulations reproduce the conformational space explored by the full simulation. As the length of the blocks increases, the sampling improves, resulting in higher covariance overlaps.

This method facilitates appraisal of a network model's predicted motions relative to increasing block lengths. The block averaged covariance overlap is plotted as a function of block size and the covariance overlap between the ENM-NMA and our gold-standard (the full-length simulation PCA) is then overlaid.^{35,40} The intersection of the block averaging curve and the ENM-NMA is the point where they have the same covariance overlap. Following this value to the x-axis gives an estimate of the simulation time required before the trajectory would predict the system's motions as well as the ENM.

2.2.2 Comparing Different Resolution ENMs—In calculating covariance overlaps of different resolution models, we noticed that the dimensionality of the system intrinsically altered the covariance overlap; physically, this is related to the fact that in higher dimensional spaces it is harder to find two vectors with large dot products. Thus, we needed to find a means to compare models with varying resolution on an equal footing.

Our first idea was to compute a Z-score, comparing the eigendecomposition from a given ENM to that of a number of randomly generated covariance matrices. However, the random matrices tended to produce very flat eigenvalue spectra, as opposed to the rapidly decaying ones produced by ENM-NMA and simulation PCA. The eigendecompositions from the random matrices gave large covariance overlaps with each other, as they would with any other eigenset (set of eigenpairs), invalidating their use as a control.

To work around the covariance overlap's sensitivity to the shape of the eigenvalue spectrum, we randomly scrambled the eigenvectors such that they were associated with different eigenvalues. We then performed a bootstrap-like analysis, generating 10,000 eigensets in this manner and computing their covariance overlap to the full trajectory. This was used to compute the mean ($\bar{\Omega}$) and standard deviation (σ) of the resulting distribution. Finally, we calculated the Z-score for the ENM's eigenset using:

$$Z = \frac{\Omega - \bar{\Omega}}{\sigma} \quad (7)$$

3 Results

3.1 Any Spring Function can Perform Well

We quantified the performance of several different spring functions and resolutions of network models by fitting to MD simulations. In order to determine the robustness of the parametrization, each spring function was fit to the entire data set (all three proteins) and to each trajectory individually. The results in Figure 1 and Supplementary Tables S1–S3 show that all spring functions considered yielded models which predict motions with similar fidelity. Figure 1A shows the results for spring functions parametrized using the β_2AR data. The covariance overlap between an ENM using the indicated spring function and the PCA of the MD simulation was plotted. Panels B and C show equivalent data for rhodopsin and CB2, respectively. In each case, the data is clustered by model resolution.

We found that after fitting, all spring functions for a particular protein and resolution produced similar overlaps. This result was quite surprising considering the scope of the literature based on defining and improving spring functions. Comparison of the various potentials revealed the basis of this similar performance; as shown in Figure 2, all of the models produce very strong interactions between each bead and the first surrounding shell of beads, while more distant interactions are roughly an order of magnitude weaker. The standard anisotropic network model, with spring contacts represented by the Heaviside step function, also follows this pattern: the fitting produces a very short cutoff such that only immediately neighboring beads are connected. A close inspection of the literature reveals this same trend of stiffening close contacts in the initial development of these methods.^{24,27,30,31,45} In fact, the original Rouse model only connected beads neighboring in sequence.⁶³ Specifically addressing ENMs, Hinsen et al. found the transition to be 4.0 \AA^{24} and Moritsugu & Smith gave virtual bonds a stiffness two orders of magnitude higher than virtual angles.³⁰ The vertical dotted line drawn at 3.8 \AA in Figure 2 represents the mean distance of the first C_α shell.

3.2 Performance is Affected by Model Resolution

In examining Figure 1, it is immediately clear that ENM performance is affected by model resolution. For example, the two-bead per residue model yielded a lower covariance overlap in all cases. The overlaps and spring parameters used are completely documented in Supplementary Tables S1–S4. The lower overlap appears to be due to three phenomena. First, a higher dimensional system attempts to predict the dynamics with higher detail, in effect giving the ENMs a tougher problem to solve. Second, side chains often have more than one rotameric minimum and thus do not obey the underlying harmonic assumption characteristic of network model normal modes. Finally, working in a higher dimensional space intrinsically reduces the dot products between vectors, since the entropic penalty for making two vectors parallel increases dramatically.

This shows up in the PCA results from the MD as well; Figure 3 shows that the block averaged covariance overlap is lower than in the one-bead per residue resolution PCA. However, the drop in ENM overlap when changing resolution outpaced the new block averaging result. This is illustrated in Figure 4, which shows how the various ENMs perform for rhodopsin. In panel A, the similarity to MD of the one-bead per residue models is plotted over the block averaging result. With standard resolution, the ENM normal modes sample approximately as well as a 400–450 ns simulation. This contrasts with panel B, which shows

the two-bead per residue result where the ENMs reproduce the dominant fluctuations with the accuracy of about 350 ns of simulation. For all spring functions tested the higher-resolution model produced modes less similar to the MD.

3.2.1 Resolution Impacts Significance of Covariance Overlap—To assess the statistical significance of the lower covariance overlaps, we used a standard measure, the Z-score. A bootstrapping-like procedure was used to generate a large number of fake data sets with the appropriate dimensionality and eigenvalue spectrum (see Section 2.2.2 for details). The Z-score analysis showed that the significance of a particular covariance overlap increased with model resolution. The aggregate Z-scores are plotted against their respective covariance overlaps in Figure 5. Here each datum is a particular spring formalism. The significance, which is plotted as a function of covariance overlap, shows that the same overlap is more difficult to achieve in the higher resolution models.

3.3 Fitting to a Single Protein is Sufficient

Our next step was to examine ENM dependence on the dataset used for parametrization. To do so, we performed a cross-validation. Specifically, we fit ENMs to each MD trajectory individually and used the resulting spring parameters to compute the motions for the other proteins. For completeness, we also tested the performance of ENMs fit to the 3 trajectories in aggregate. The results, shown in Figure 6, indicate very little difference in performance regardless of which data set was used for fitting.

The lack of improvement upon individualized fitting stems from homogeneity in the spring parameters. As mentioned in Section 3.1, all functional forms tend to favor stiffening the first shell interactions much more than distal contacts. In fitting Eq. (2) this was realized by a shortened cutoff distance, which can be seen in Table 1. In all cases, using the new parameters, increased the covariance overlaps to their respective simulation PCA results.

Although this result may be due in part to the similarity of the three proteins used for fitting (all class A GPCRs), there is some level of generality in fitting network model parameters — at least among the proteins tested. In order to further validate our parameters, we compared them to simulations of two additional GPCRs not used for fitting: A2A (2.2 μ s in aggregate) and opsin (472 ns). The opsin case is particularly challenging, since the structure more closely resembles an active GPCR structure, while the others clearly model the inactive forms of the protein. Table 1 lists their covariance overlaps using the typical 15 Å cutoff (before fitting) and the cutoffs we obtained via parameterization. We found that the covariance overlaps for both proteins behave much as the ones used in the fitting: the fit parameters perform significantly better than the standard ones in both the one- and two-beads per residue models, with covariance overlaps comparable to those for the other proteins. The covariance overlaps of the one-bead per two-residues model were similar before and after fitting. This is unsurprising, because fitting only shortened the cutoff by ≈ 1 Å.

3.4 Does the Quality of the Fitting Results Depend on the Trajectory Length?

Our previous work suggested that fitting to short MD trajectories was a flawed strategy because, in addition to larger statistical errors, short trajectories also differ systematically from longer ones. Specifically, the eigenvalues computed from PCA on short trajectories decay more slowly. Physically, this means that short trajectories systematically underestimate the importance of the largest motions in the system. In that manuscript (ref. 35), we argued that fitting against short trajectories would recapitulate these problems, and thus should be avoided.

In the present study, we have shown that the overall quality of the results, as measured by covariance overlap, does not seem to depend sensitively on the data used in the fitting. To clarify the apparent contradiction, we explicitly tested ENM performance while varying the length of the trajectory used for fitting.

To do this, the first microsecond of each long trajectory was split into contiguous pieces, yielding sets of 25, 50, 100, 250, and 500 ns trajectories. These shorter trajectories were then used to parametrize the various ENM formalisms in the same manner described in Section 2.2. The resulting parameters were then used to perform ENM calculations that were compared to the entire trajectory: 1.02 μ s for β_2 AR, 1.6 μ s for rhodopsin, and 1.9 μ s for CB2. Surprisingly, ENMs parametrized against short trajectories perform nearly as well as those fit to long trajectories, as shown in Figure 7.

Following the results of parametrizing with short trajectories we wanted to further explore why ENM-NMA performance was not affected by the trajectory length used for fitting. We previously noted that the power spectra from PCA applied to trajectories varied systematically with the length of the trajectory, as seen in Figure 8.³⁵ Panel A shows that as the simulation time increases the fractional contribution made by the lowest mode increases. Note that this is not due to larger motions by the protein itself — the RMSD to the starting structure plateaus within the first 100 ns — but rather due to the dynamics, as longer time-scale motions are gradually accessed by the trajectory.

From our previous data we predicted an improvement of ENM power spectra after fitting. However, in the present work, we found that the difference between ENM power spectra after fitting a short or a long trajectory is negligible. An example is shown in panel B of Figure 8, using the HCA formalism. Here, the power spectra averaged over the entire ensemble were plotted for each of the trajectory lengths tested. This includes ENMs fit to each of the shorter trajectories along with the those fit to the full trajectory. The error bars represent one standard deviation from the ensemble average of shorter trajectories. The fractional contribution of a particular mode to the overall motion is plotted for the first 10 eigenvalues in the ENM-NMA, with modes increasing along the horizontal axis.

Interestingly, the ENM-NMA power spectrum, shown in panel B, is actually slightly “softer” (the lowest modes contribute less to the overall motion and the contribution decays more slowly) after fitting to longer trajectories. This trend was seen in all distance-dependent spring functions. However, the change in fractional contribution of the most dominant mode is much smaller than the change seen in the PCA of increasingly longer simulations. By contrast, the power spectrum of a simulation’s principal components (Figure 8A) changes drastically with increasing time. Fitting improved the ENM-NMA power spectrum, but not enough to keep up with the changes in fractional contribution of the simulation’s principal components. Thus, it appears that this flaw is fundamental to the ENM formalism. One possible explanation is that the lowest frequency modes in the MD trajectories are transitions between similar substates, as opposed to fluctuations about a single minimum. These motions are intrinsically absent in an ENM-NMA formalism, and this absence could account for the lack of improvement in their performance.

3.4.1 Low Frequency PCA Modes Show Transitions between Substates—In order to assess whether transitions appear in the simulation we histogrammed the projection of our rhodopsin trajectory onto the lowest principal components (also called the right singular vectors).^{35,64} The peaks in such a histogram indicate the highly populated conformational substates along a given mode. Multiple well-distinguished peaks suggest that there is more than one distinct state along that vector, whereas unimodal histograms indicate fluctuations about a single state. The latter case is well-modeled by a harmonic

approximation, while the former is not. As can be seen in Figure 9, time series for the slowest modes of a one-bead per residue PCA show evidence of structural transitions — the motions along these modes have multiple peaks. These low-mode transitions indicate the sampling of separate substates within the most concerted motions in the trajectory. These are precisely the kinds of anharmonic motions that ENMs are unable to predict; even if they accurately detect the direction of motion, they will not be “aware” of the second energy well, and will thus systematically underestimate the amplitude of the motion along the mode.

3.5 ENMs are Robust to Changes in Parameters

In order to systematically explore the dependence of ENM performance on the chosen parameters, we explicitly tessellated the parameter space for one model (HCA) using the one-bead per residue resolution. We found the two most important parameters were the cutoff distance, R_c , and the exponent weighting, d . Figure 10 demonstrates that the covariance overlap is quite robust to systematic variation of these parameters. In Figure 10 the value of the exponential (“ d ” in Eq. (4)) is plotted on the vertical axis and cutoff distance (“ R_c ” in Eq. (4)) on the horizontal axis, with the resulting covariance overlap indicated by the heat map’s intensity. This illustrates that a broad range of parameters ($d = 7-14$ and $R_c = 1-5$) yielded similar overlap with the simulation’s principal components. This degeneracy in parameters partially explains ENM’s robustness to the variations we tested, revealing how many spring functions and even shorter trajectory fits arrive at such similar results. Interestingly, at $R_c < 3.8 \text{ \AA}$ nearly all connections are represented by the exponentially decaying spring. Thus, even changing the spring function does not significantly impact the overlap.

4 Discussion

Through the use of long trajectories and an eigenvalue-sensitive metric we were able to quantitatively show and clarify the underlying generality of network model parameters. We found that this robustness makes network models relatively easy to parametrize and that choice of spring function makes little difference on performance. Interestingly, we also found essentially no dependence on the length of trajectory used for parametrization. However, this was largely because the network models were not able to accurately reproduce the eigenvalue spectrum of a simulation. Some other variation in the formalism that is able to better capture this power spectrum might well benefit from being fit only to long trajectories.

4.1 ENMs Rely on Short-Range Contacts

With the numerous formalisms in print, we set out to measure the accuracy of each ENM with our set of long simulations. This dataset has the advantage of reproducing fluctuations on a much longer time scale than was used for earlier ENM parametrizations, so there is a higher probability of capturing the slow motions predicted by a network model.

As shown in Figure 1 and Figure 4, the choice of spring function has little bearing on the performance of a network model. Figure 2 brings the reason to light: all spring functions tend to favor much stiffer springs in the first C_α shell. Interestingly, Tozzini reported this same trend in a recent review of coarse-grained models for protein motion.⁶⁵ This evidence argues against any single functional form and elucidates their most striking commonality, reliance on the short-range contacts.

4.2 Higher Resolution Models Show Lower Covariance Overlap

Upon fitting higher resolution models the covariance overlap decreased regardless of spring function. This decrease was more rapid than that found in block-averaged PCA results. In order to characterize this dependency on resolution, we implemented a Z-score test (using decoupled eigenpairs to calculate an average overlap as described in Section 2.2.2). Although the absolute covariance overlap dropped, the statistical significance generally improved. Moreover, the simulation time required to produce equivalent covariance overlaps is longer for the higher resolution models (see for example Figure 3), suggesting that the higher resolution models do perform well.

4.3 Parameters are Independent from Protein Parametrized

Although our parametrization set consisted of only three simulations, each of a class A GPCR, the result of cross-validating the parametrization of each protein is important. The high overlap when cross-validating ENM normal modes with various proteins reveals two key details. The first is that there is some level of generality within the proteins we tested, and the results of our parametrization can be used in the future to study GPCRs. To this effect we tested our shorter Heaviside cutoffs on simulations of two other GPCRs, opsin and A2A. The results in Table 1 show that the fit parameters perform better than the defaults for both proteins, although the change is insignificant for the one-bead per two-residues resolution, and that the performance for these proteins is comparable to that for the ones used in the fitting. Thus, we conclude that ENMs are robust to variations in spring parameters, as illustrated by Figure 6 and Figure 10. That said, we believe it is important in the future to extend this test set to other protein architectures, particularly globular proteins.

Our results explain the ease of finding high overlaps and, considered together with the rigidity of first shell contacts, show that any ENM can be quickly parametrized “by hand” by stiffening those contacts. Another method would be to discretely represent the first shell separately from all other contacts, as in Eq. (5). This is in fact used as part of the REACH³⁰ and edENM³³ formalisms, but the separation of more distal contacts presented in these methods was not found to be critical to ENM performance.

Taken together, these two pieces of evidence demonstrate the robustness of network models. In light of these results, we plotted the parameter space of the HCA formalism. Figure 10 illustrates the broad range of ENM parameters that yield high covariance overlaps. This demonstrates ENMs’ insensitivity to choice of parameters and functional form.

4.4 ENMs are Insensitive to Trajectory Length used for Parametrization

The results shown in Figure 7 were surprising. Our previous study showed that the fractional contribution of motions in MD simulations had a strong time dependence and that network models better reproduced this curve after fitting.³⁵ While it is clear from Figure 8A that the power spectrum of a simulation PCA result changes shape with trajectory length, closer inspection of panel B shows that the power spectrum of an ENM-NMA result still underestimates the contribution of the first few modes. Also, Figure 8 showed that the power spectra of ENM-NMAs fit to different length trajectories varied little. Therefore, it appears that the parametrization improves overlap in a manner independent from the trajectory actually fit. This revealed a flaw inherent in all network models we parametrized: they were unable to correctly predict the relative importance of the lowest-frequency modes of long simulations. One possible explanation is shown in Figure 9. The first three principal components of the simulations contain transitions (data shown for rhodopsin). ENM-NMA, which uses a harmonic representation, is unable to predict such motions. In the future, an ENM that could accurately describe the power spectrum of a long simulation would likely

benefit from fitting to only long trajectories. In this respect, the “double-native” G model proposed by Zuckerman and co-workers is quite intriguing.⁶⁶

In addition, incorporating more simulations of diverse molecules into this study without relaxing our restriction of trajectory length may improve our interpretation. GPCR transmembrane regions are quite rigid, so including other proteins would further elucidate the behavior of eigenvalue contributions. Including different classes of proteins would also be significant as our results showed a very general trend in fitting — namely stiffening the first shell contacts.

5 Conclusions

We evaluated many formalisms present in the elastic network model literature using a set of long (each $> 1 \mu\text{s}$) molecular dynamics simulations. The principal components of these simulations were the basis for comparison, evaluation, and fitting of network models. Fitting was preformed using an eigenvalue-sensitive metric, the covariance overlap, in order to capture the relative amplitude of motions present in both systems. Somewhat surprisingly, all of the formalisms showed roughly equivalent performance regardless of the data used for parametrization. This result could be viewed as disappointing — sophisticated efforts to improve the models are largely unsuccessful — or as informative due to the similarities in parameters between functional forms. Our results showed that fitting caused all formalisms to assign strong interactions between first shell contacts and much weaker interactions for longer range contacts.

This led us to tessellate the parameter space in several models. We found that the spring parameters were fairly robust with broad ranges yielding high covariance overlaps. Therefore, we conclude that while significant improvement can come from fitting the spring parameters, the formalism complexity has little bearing on the actual accuracy of the model. In addition, a high performance parametrization can likely be obtained easily by assigning stiff connections between C_a 's neighboring in sequence and much weaker connections to more distal contacts.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors would like to thank Joshua Horn, Dejun Lin, and Jessi Leioatts for their critical review and suggestions to the manuscript. We also thank Edward Lyman and Ji Young Lee for their generous donation of the A2A receptor simulations, the University of Rochester Center for Integrated Research Computing for the resources to run the opsin simulation, and Michael Pitman and the IBM T. J. Watson Research Center for the resources required to run the simulations of $\beta_2\text{AR}$, CB2, and rhodopsin. This work was supported by NIH project number 1R01GM095496 and NL is supported by the Institutional Ruth L. Kirschstein National Research Service Award, grant number GM068411.

References

1. Shaw, DE.; Dror, RO.; Salmon, JK.; Grossman, JP.; Mackenzie, KM.; Bank, JA.; Young, C.; Deneroff, MM.; Batson, B.; Bowers, KJ.; Chow, E.; Eastwood, MP.; Ierardi, DJ.; Klepeis, JL.; Kuskin, JS.; Larson, RH.; Lindorff-Larsen, K.; Maragakis, P.; Moraes, MA.; Piana, S.; Shan, Y.; Towles, B. Millisecond-scale molecular dynamics simulations on Anton. Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis; 2009.
2. 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, Newns 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. *IBM Syst J*. 2001; 40:310.
3. Shirts M, Pande VS. *Science*. 2000; 290:1903–1904. [PubMed: 17742054]
 4. Bahar I, Lezon TR, Bakan A, Shrivastava IH. *Chem Rev*. 2010; 110:1463–1497. [PubMed: 19785456]
 5. Jernigan RL, Demirel MC, Bahar I. *Int J Quantum Chem*. 1999; 75:301–312.
 6. Tama F, Valle M, Frank J, Brooks CL. *Proc Natl Acad Sci U S A*. 2003; 100:9319–9323. [PubMed: 12878726]
 7. Wang Y, Rader AJ, Bahar I, Jernigan RL. *J Struct Biol*. 2004; 147:302–314. [PubMed: 15450299]
 8. Zheng W, Brooks BR. *Biophys J*. 2005; 89:167–178. [PubMed: 15879477]
 9. Isin B, Rader AJ, Dhiman HK, Klein-Seetharaman J, Bahar I. *Proteins: Struct, Funct, Bioinf*. 2006; 65:970–983.
 10. Yang L, Song G, Jernigan RL. *Biophys J*. 2007; 93:920–929. [PubMed: 17483178]
 11. Zheng W, Brooks BR, Thirumalai D. *Biophys J*. 2007; 93:2289–2299. [PubMed: 17557788]
 12. Niv MY, Filizola M. *Proteins: Struct, Funct, Bioinf*. 2008; 71:575–586.
 13. Ahmed A, Villinger S, Gohlke H. *Proteins: Struct, Funct, Bioinf*. 2010; 78:3341–3352.
 14. Chennubhotla C, Rader AJ, Yang LW, Bahar I. *Phys Biol*. 2005; 2:S173–S180. [PubMed: 16280623]
 15. Rueda M, Chacóshon P, Orozco M. *Structure*. 2007; 15:565–575. [PubMed: 17502102]
 16. Yang LW, Eyal E, Chennubhotla C, Jee J, Gronenborn AM, Bahar I. *Structure*. 2007; 15:741–749. [PubMed: 17562320]
 17. Delarue M, Dumas P. *Proc Natl Acad Sci U S A*. 2004; 101:6957–6962. [PubMed: 15096585]
 18. Tama F, Miyashita O, Brooks CL. *J Mol Biol*. 2004; 337:985–999. [PubMed: 15033365]
 19. Zheng W. *Biophys J*. 2011; 100:478–488. [PubMed: 21244844]
 20. Grossfield A, Zuckerman DM. *Annu Rep Comput Chem*. 2009; 5:23–48. [PubMed: 20454547]
 21. Woodcock HL, Zheng W, Ghysels A, Shao Y, Kong J, Brooks BR. *J Chem Phys*. 2008; 129:214109. [PubMed: 19063546]
 22. Ma J. *Structure*. 2005; 13:373–380. [PubMed: 15766538]
 23. Atilgan AR, Durell SR, Jernigan RL, Demirel MC, Keskin O, Bahar I. *Biophys J*. 2001; 80:505–515. [PubMed: 11159421]
 24. Hinsin K, Petrescu AJ, Dellerue S, Bellissent-Funel MC, Kneller GR. *Chem Phys*. 2000; 261:25–37.
 25. Kovacs JA, Chacón P, Abagyan R. *Proteins: Struct, Funct, Bioinf*. 2004; 56:661–668.
 26. Micheletti C, Carloni P, Maritan A. *Proteins: Struct, Funct, Bioinf*. 2004; 55:635–645.
 27. Jeong JI, Jang Y, Kim MK. *J Mol Graphics Modell*. 2006; 24:296–306.
 28. Kondrashov DA, Cui Q, Phillips GN. *Biophys J*. 2006; 91:2760–2767. [PubMed: 16891367]
 29. Song G, Jernigan RL. *Proteins: Struct, Funct, Bioinf*. 2006; 63:197–209.
 30. Moritsugu K, Smith JC. *Biophys J*. 2007; 93:3460–3469. [PubMed: 17693469]
 31. Lyman E, Pfaendtner J, Voth GA. *Biophys J*. 2008; 95:4183–4192. [PubMed: 18658214]
 32. Yang L, Song G, Jernigan RL. *Proc Natl Acad Sci U S A*. 2009; 106:12347–12352. [PubMed: 19617554]
 33. Orellana L, Rueda M, Ferrer-Costa C, Lopez-Blanco JR, Chacón P, Orozco M. *J Chem Theory Comput*. 2010; 6:2910–2923.
 34. Yang LW. *Biophys J*. 2011; 100:1784–1793. [PubMed: 21463592]
 35. Romo TD, Grossfield A. *Proteins: Struct, Funct, Bioinf*. 2011; 79:23–34.
 36. Doruker P, Atilgan AR, Bahar I. *Proteins: Struct, Funct, Bioinf*. 2000; 40:512–524.
 37. Brooks BR, Janežič D, Karplus M. *J Comput Chem*. 1995; 16:1522–1542.
 38. Hess B. *Phys Rev E Stat, Nonlinear Soft Matter Phys*. 2002; 65:031910.

39. Lyman E, Zuckerman DM. *J Phys Chem B*. 2007; 111:12876–12882. [PubMed: 17935314]
40. Romo TD, Grossfield A. *J Chem Theory Comput*. 2011; 7:2464–2472.
41. Liu L, Gronenborn AM, Bahar I. *Proteins: Struct, Funct, Bioinf*. 2011; 80:616–625.
42. Tirion S. *Phys Rev Lett*. 1996; 77:1905–1908. [PubMed: 10063201]
43. Bahar I, Atilgan AR, Erman B. *Folding Des*. 1997; 2:173–181.
44. Tama F, Gadea FX, Marques O, Sanejouand YH. *Proteins: Struct, Funct, Bioinf*. 2000; 41:1–7.
45. Ming D, Wall ME. *Phys Rev Lett*. 2005; 95:198103. [PubMed: 16384030]
46. Riccardi D, Cui Q, Phillips GN. *Biophys J*. 2010; 99:2616–2625. [PubMed: 20959103]
47. Grossfield A, Feller SE, Pitman MC. *Proteins: Struct, Funct, Bioinf*. 2007; 67:31–40.
48. Hinsen K. *Proteins: Struct, Funct, Bioinf*. 1998; 33:417–429.
49. Eyal E, Yang LW, Bahar I. *Bioinformatics*. 2006; 22:2619–2627. [PubMed: 16928735]
50. Romo TD, Grossfield A. *Conf Proc IEEE Eng Med Biol Soc*. 2009; 2009:2332–2335. [PubMed: 19965179]
51. Romo, TD.; Grossfield, A. LOOS: Lightweight Object Oriented Structure analysis. Grossfield Lab, University of Rochester Medical School; Rochester, NY: <http://loos.sourceforge.net>
52. Bahar I, Jernigan RL. *J Mol Biol*. 1997; 266:195–214. [PubMed: 9054980]
53. Lau P-W, Grossfield A, Feller SE, Pitman MC, Brown MF. *J Mol Biol*. 2007; 372:906–917. [PubMed: 17719606]
54. Grossfield A, Pitman MC, Feller SE, Soubias O, Gawrisch K. *J Mol Biol*. 2008; 381:478–486. [PubMed: 18585736]
55. Romo TD, Grossfield A, Pitman MC. *Biophys J*. 2010; 98:76–84. [PubMed: 20074514]
56. Hurst DP, Grossfield A, Lynch DL, Feller S, Romo TD, Gawrisch K, Pitman MC, Reggio PH. *J Biol Chem*. 2010; 285:17954–17964. [PubMed: 20220143]
57. Feller SE, Gawrisch K, MacKerell AD. *J Am Chem Soc*. 2002; 124:318–326. [PubMed: 11782184]
58. Klauda JB, Brooks BR, MacKerell AD, Venable RM, Pastor RW. *J Phys Chem B*. 2005; 109:5300–5311. [PubMed: 16863197]
59. MacKerell AD Jr, Bashford D, Bellott M, Dunbrack R Jr, Evanseck J, Field M, Fischer S, Gao J, Guo H, Ha S, Joseph-McCarthy D, Kuchnir L, Kuczera K, Lau F, Mattos C, Michnick S, Ngo T, Nguyen D, Prodhom B, Reiher W III, Roux B, Schlenkrich M, Smith J, Stote R, Straub J, Watanabe M, Wiorkiewicz-Kuczera J, Yin D, Karplus M. *J Phys Chem B*. 1998; 102:3586–3616.
60. Lee JY, Lyman E. *Biophys J*. 2012; 102:2114–2120. [PubMed: 22824275]
61. Faraldo-Gómez JD, Forrest LR, Baaden M, Bond PJ, Domene C, Patargias G, Cuthbertson J, Sansom MSP. *Proteins: Struct, Funct, Bioinf*. 2004; 57:783–791.
62. Nelder JAMR. *Comput J*. 1964; 7:308–313.
63. Rouse PE. *J Chem Phys*. 1953; 21:1272–1280.
64. Romo TD, Clarage JB, Sorensen DC, Phillips GN. *Proteins: Struct, Funct, Bioinf*. 1995; 22:311–321.
65. Tozzini V. *Q Rev Biophys*. 2010; 43:333–371. [PubMed: 20707940]
66. Zuckerman DM. *J Phys Chem B*. 2004; 108:5127–5137.

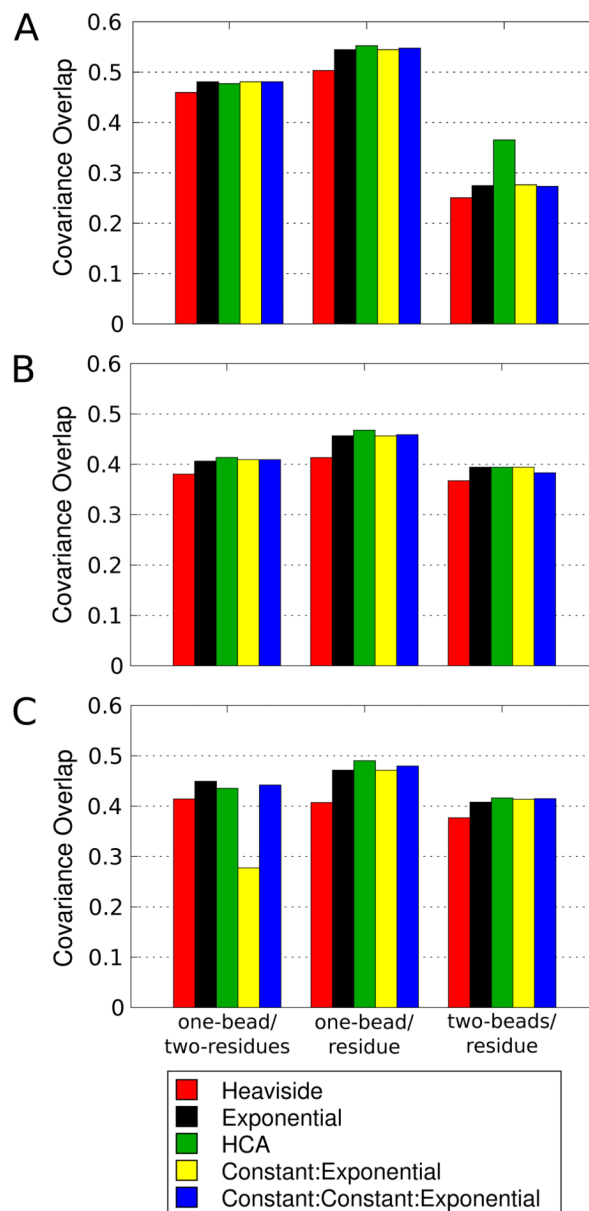


Figure 1.

Covariance overlap, Eq. (6), between an ENM result and the full length simulation PCA for representative spring functions clustered by model resolution. “Heaviside” refers to the standard distance cutoff method described by Eq. (2). “Exponential” form is defined by Eq. (3), “HCA” is defined by Eq. (4), and “constant” implies a scalable constant value, for instance k_1 in Eq. (5). Spring names separated by colons (:) represent functions where explicit bonding was used. Here, C_a ’s neighboring in sequence are bonded (n to $n+1$). When the two-bead per residue resolution is used each C_a bead is bonded to its own C_{CoM} bead as well. The hierarchy of explicit levels of bonding takes the form (n to $n+1$: n to $n+2$: ... : n to non-bonded). All values reported were fit to the full simulation using a Nelder-Mead Simplex⁶² optimization to maximize the covariance overlap. Data is shown for (A) β_2AR , (B) rhodopsin, and (C) CB2.

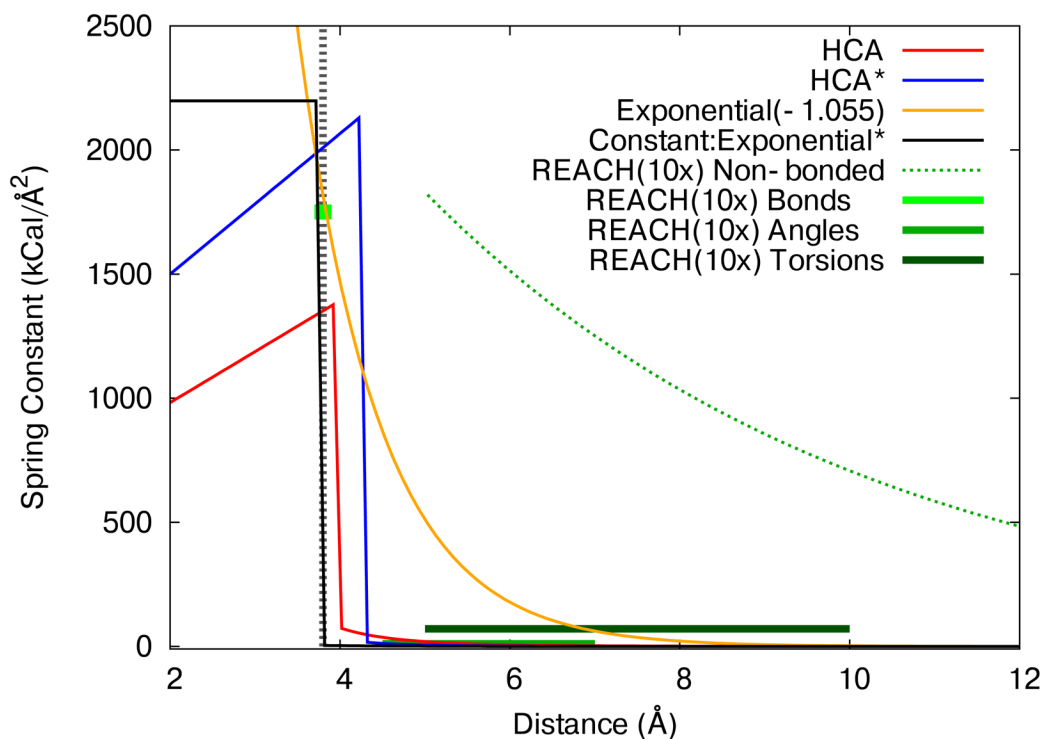


Figure 2.

Spring stiffness as a function of distance. Several formalisms are shown. “HCA” refers to the function developed by Hinsen et al.²⁴ and defined here by Eq. (4). “HCA*” is the same function, but uses values obtained from our parametrization.³⁵ “Exponential” is defined by Eq. (3), using $a = 1.055$, which gave the highest covariance overlap in our simultaneous fit of all three trajectories (see Supplementary Table S4) when using the standard one-bead per residue resolution. “Constant: Exponential” refers to Eq. (5), after fitting (parameters in Supplementary Table S4). Lastly, the parameters from the realistic extension algorithm via covariance Hessian (REACH)³⁰ method are shown. Here, springs are split into four hierarchical levels; virtual bonds, angles, and torsions and a non-bonded term. The parameters reported in that study were multiplied by a factor of 10 for easier visualization. The vertical dotted line drawn at $x = 3.8$ Å represents mean C_α distance.

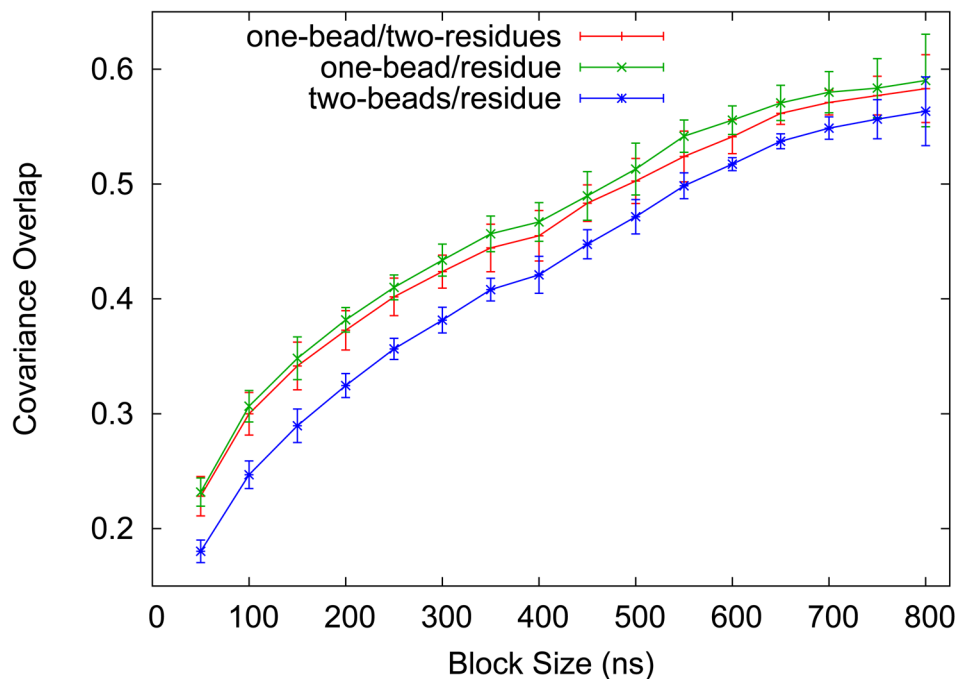


Figure 3. Block covariance overlap of rhodopsin for all three resolutions. The overlap between shorter contiguous simulation blocks are plotted against the full length trajectory with block length increasing along the x-axis. Both the higher and lower resolution models (blue and red curves, respectively) converge more slowly than the one-bead per residue resolution (green curve).

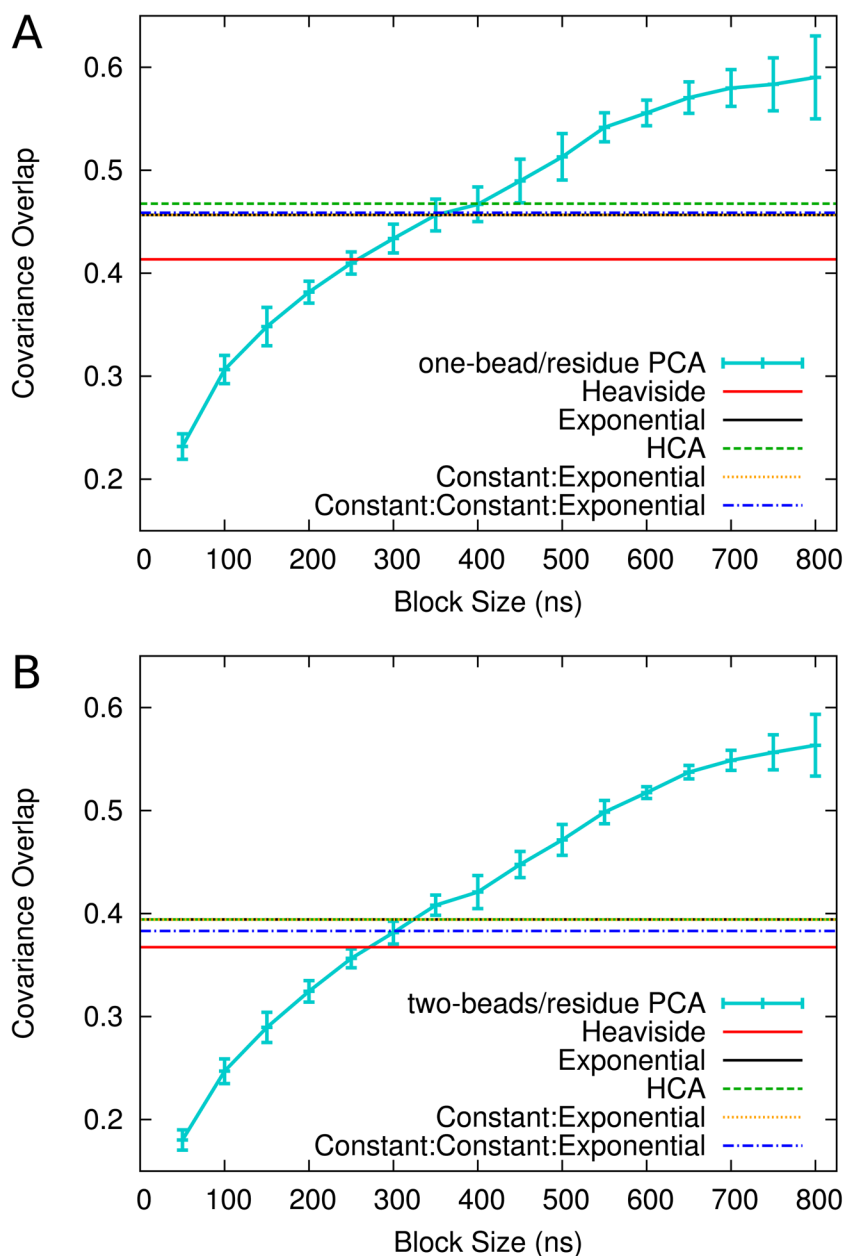


Figure 4. Rhodopsin simulation block covariance overlap plotted with ENM covariance overlaps. As in Figure 3, overlap is computed against the full trajectory. Here, ENM covariance overlaps are overlaid using the same coloring and naming convention as Figure 1. (A) Results of the standard one-bead per residue resolution, as plotted in Figure 1B, middle column. After fitting, all spring functions besides the simple distance cutoff achieved a similar overlap (0.47–0.49) which is equivalent to the block averaging result for a 375–450 ns simulation. (B) Two-bead per residue resolution (described in Section 2.1.1), as plotted in Figure 1B, right column. Here the blocked covariance overlap converges more slowly, but the ENM overlap is also lower than the standard resolution model. The covariance overlaps are again very similar between ENMs, scoring about 0.41, which was comparable to the average

overlap of 350 ns trajectory blocks. Again, the distance cutoff method performed slightly worse than the more sophisticated spring functions.

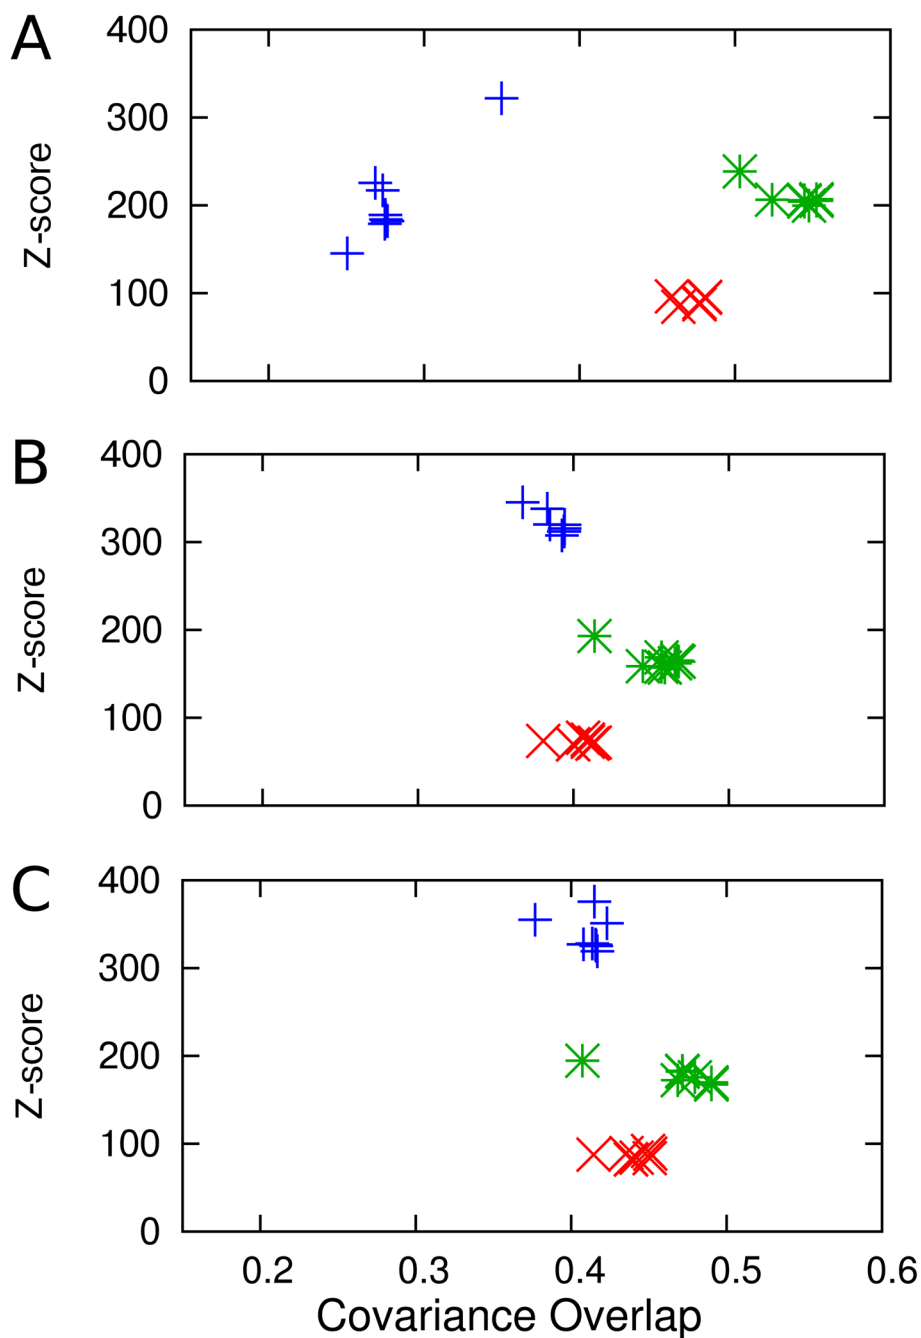


Figure 5.

Significance of covariance overlaps are dependent on model resolution. Z-scores are plotted by their covariance overlap with the full simulation PCA. We first computed the covariance overlap between ENM-NMA and the simulation PCA. After this the eigenpairs were scrambled (as described in Section 2.2.2). The new eigenset was then compared to the simulation PCA. This was done for 10,000 iterations and the mean and standard deviation of the resulting covariance overlaps was used to calculate a Z-score for the original result. Data points in red used the one-bead per two-residues resolution, green points represent one-bead per residue resolution, and blue points represent two-beads per residue resolution. Data is shown for all three proteins studied: (A) β_2 AR, (B) rhodopsin, and (C) CB2.

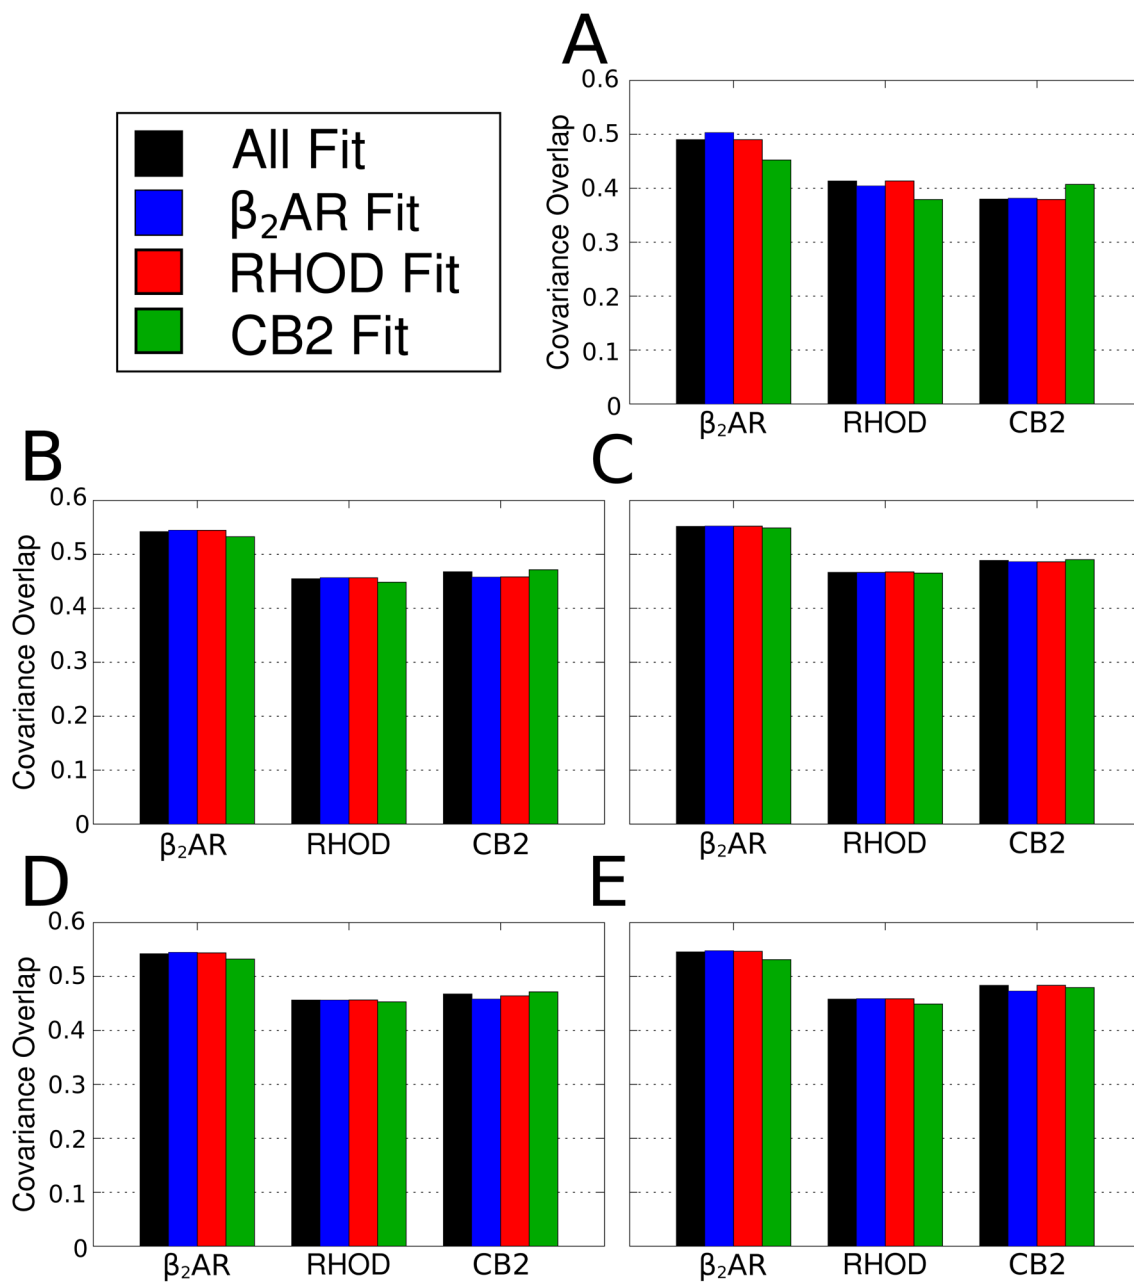


Figure 6. Simultaneous fitting yields similar accuracy as individual fitting. Here data is clustered by the protein tested. The four bars in each cluster represent the different proteins which are used in parametrization. First, all three proteins are fit simultaneously (black), then individual fits to β_2 AR (blue), rhodopsin (red), and CB2 (green) are conducted. In general, testing with the protein used for parametrization did not achieve a significantly higher covariance overlap. Simultaneous fitting also produced ENMs that performed almost as well as individualized fits. The naming convention from Figure 1 is repeated here. Results of (A) the standard distance cutoff ENM defined by the Heaviside function, Eq. (2), (B) the exponential function described by Eq. (3), (C) the HCA function defined in Eq. (4), (D) the

spring function explicitly distinguishing “bonded” terms, Eq. (5), (E) the spring function distinguishing both “bonded” and “angle” terms from “non-bonded” terms, Eq. (S8).

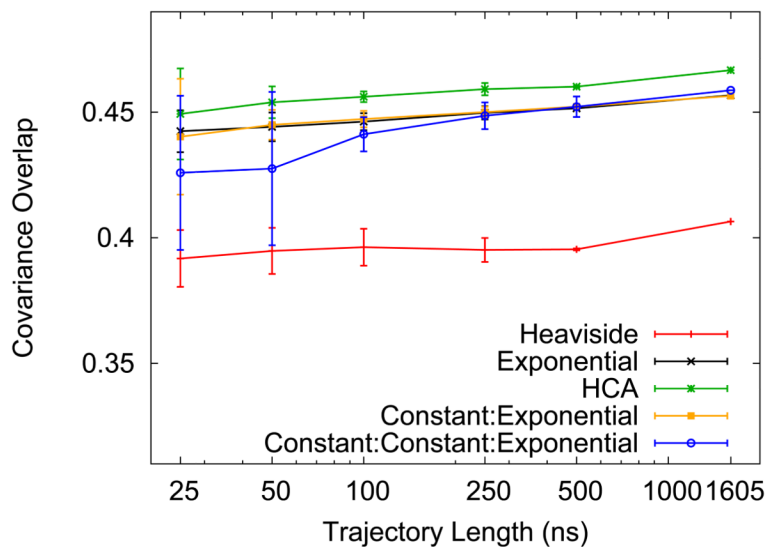


Figure 7. ENM performance does not depend on the length of trajectory used for parametrization. The first microsecond of each MD Trajectory was broken into short contiguous pieces (of 25, 50, 100, 250, and 500 ns) and then used to fit the various ENM spring formalisms. The average covariance overlap across an ensemble (\pm) one standard deviation is plotted as a function of the trajectory length used to parametrize the model (note log-scaled x-axis). Here data is shown for rhodopsin using the typical one-bead per residue resolution. Curves show different spring functions, labelled as described in Figure 1. ENMs fit to longer trajectories do not have significantly higher covariance overlaps.

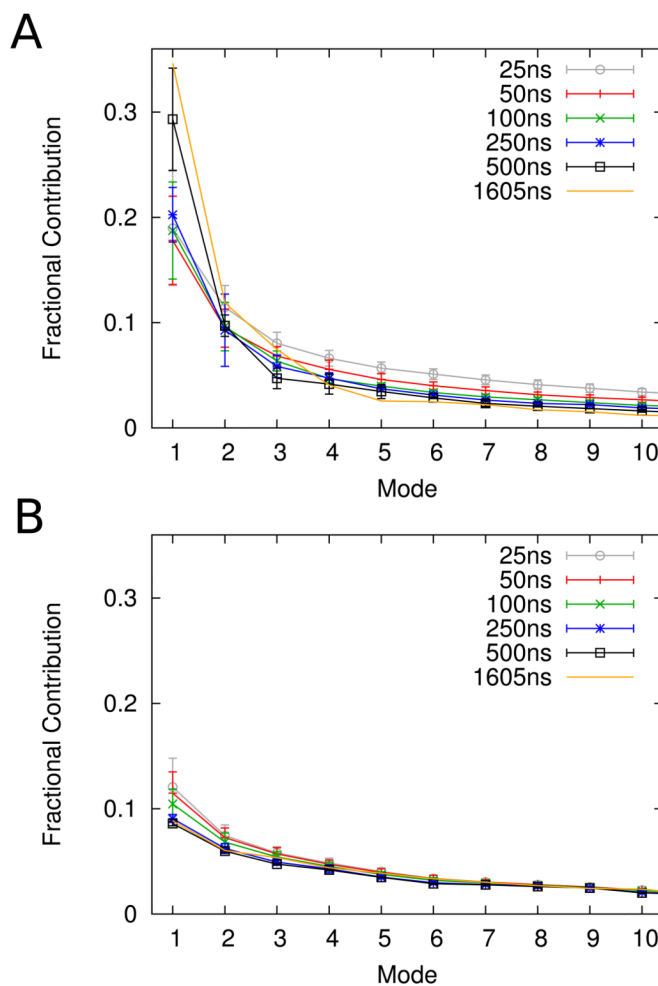


Figure 8.

ENMs fail to accurately reproduce the significance of the lowest modes of motion. The fractional contribution of the 10 most significant modes (normalized contribution of the 10 lowest frequency eigenvalues) are plotted for (A) the PCA of the simulation and (B) the HCA formalism using rhodopsin. Data is compiled by averaging the eigenvalues across each ensemble of short simulations, for a given trajectory length. Error bars indicate one standard deviation in our ensemble of shorter trajectories. (A) The fractional contribution of the lowest mode increases with longer simulations. The spectrum from PCA of the full simulation is shown in orange. The first two modes contribute more to the variance as simulation time increases. (B) Network models fail to reproduce the significance of the lowest, most concerted motions. The key denotes the MD trajectory length used to parametrize each ENM. Data from the HCA formalism²⁴ Eq. (4) is shown. Similar results were obtained with other functional forms. In all trajectory lengths used for fitting, the ENM under represents the importance of the lowest modes. In contrast to the molecular dynamics PCA, fitting longer trajectories appears to decrease the significance of these lowest modes.

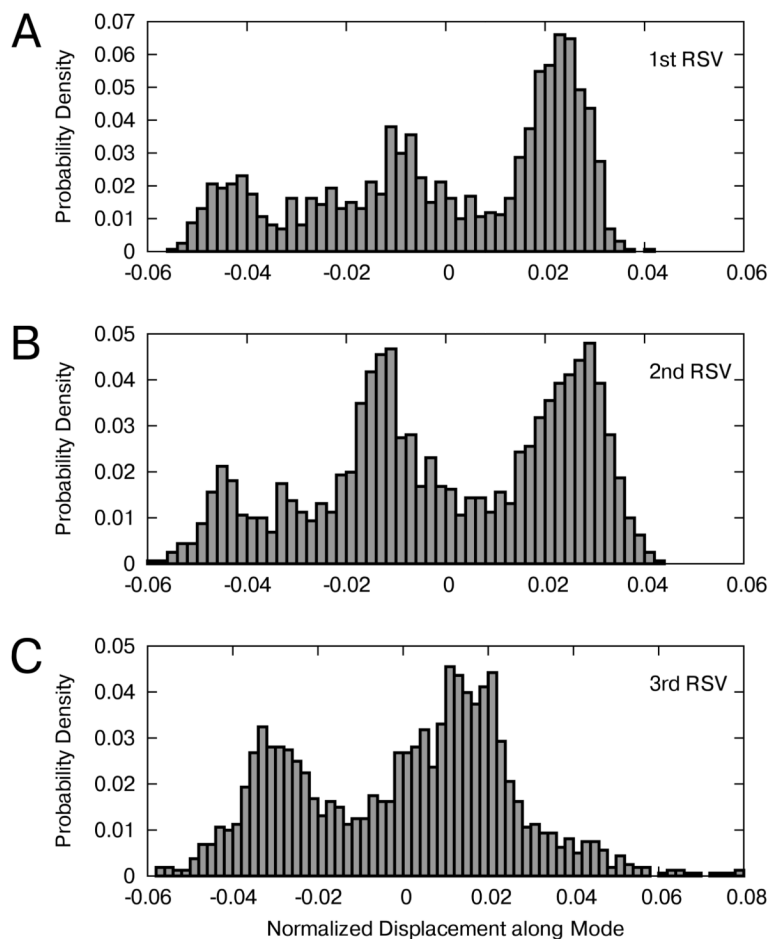


Figure 9. Histograms of right singular vectors (projection of motions onto the eigenvectors) show multiple states in the simulation's principal components. These histograms show the normalized displacement along a given principal component and were constructed using the PCA of C_{α} 's from our rhodopsin simulation. Motions along the (A) first, (B) second, and (C) third components are shown. In all three plots multiple peaks are present.




Figure 10.

A broad range of parameters are capable of yielding high overlap with simulation. The parameters d and R_c of Eq. (4) were systematically varied and compared to the PCA of the full length simulation of rhodopsin, using the covariance overlap metric, Eq. (6). Data shown used the standard one-bead per residue resolution. This revealed a large area of high overlap (from roughly $d = 6$ –14 and $R_c < 5$ Å).

Table 1

Covariance overlaps using the standard Heaviside function (Eq. (2)) after parametrization against β_2 AR, rhodopsin, and CB2 in aggregate. The data is shown for all three resolutions tested, one-bead/two-residues (1/2), one-bead/residue (1/1), and two-beads/residue (2/1). A2A and opsin were not parametrized. Covariance overlaps between these ENM-NMAs and simulation PCA is reported first using a standard 15 Å cutoff distance and then after using the parameters obtained by fitting the other three proteins (cutoff distances indicated).

Protein	Covariance Overlap		
	1/2	1/1	2/1
β_2 AR	0.442	0.490	0.251
rhodopsin	0.380	0.413	0.358
CB2	0.413	0.380	0.374
Before Fit	$R_c = 15 \text{ \AA}$		
A2A	0.361	0.305	0.305
opsin	0.330	0.273	0.222
After Fit	$R_c = 13.93 \text{ \AA}$	$R_c = 8.636 \text{ \AA}$	$R_c = 8.628 \text{ \AA}$
A2A	0.362	0.410	0.410
opsin	0.327	0.363	0.296