J Mol Model
DOI 10.1007/s00894-009-0504-3

ORIGINAL PAPER

Characterizing rare-event property distributions via replicate molecular dynamics simulations of proteins

Ranjani Krishnan · Emily B. Walton · Krystyn J. Van Vliet

Received: 16 February 2009 / Accepted: 9 March 2009

Abstract As computational resources increase, molecular dynamics simulations of biomolecules are becoming an increasingly informative complement to experimental studies. In particular, it has now become feasible to use multiple initial molecular configurations to generate an ensemble of replicate production-run simulations that allows for more complete characterization of rare events such as ligand-receptor unbinding. However, there are currently no explicit guidelines for selecting an ensemble of initial configurations for replicate simulations. Here, we use clustering analysis and steered molecular dynamics simulations to demonstrate that the configurational changes accessible in molecular dynamics simulations of biomolecules do not necessarily correlate with observed rare-event properties. This informs selection of a representative set of initial configurations. We also employ statistical analysis to identify the minimum number of replicate simulations required to sufficiently sample a given biomolecular property distribution. Together, these results suggest a general procedure for generating an ensemble of replicate simulations that will maximize accurate characterization of rare-event property distributions in biomolecules.

Keywords Clustering analysis · Energy landscape · Ligand-receptor · Property distribution · Replicate simulations · Steered molecular dynamics

Introduction

Molecular dynamics (MD) simulations have been utilized to model and gain mechanistic understanding of protein behavior for over 30 years [1]. The timescale of MD simulations has historically been much shorter than that of experiments, but with increased computational capacity it has become possible to perform both longer simulations and more simulations. For example, it is now possible to attain simulation durations for solvated proteins on the timescale of physical events (micro-second timescales) [2, 3]. It is also now possible to conduct replicate simulations that explore the variation as well as the magnitude of biomolecular properties [4, 5].

MD studies primarily utilize a single, long-timescale trajectory to draw conclusions about time-averaged biomolecular properties [2, 6]. For investigation of time-averaged properties, it is clear that a single long trajectory is desirable and informative. However, MD simulations of proteins are also used to investigate so-called "rare events" which occur infrequently in a simulation trajectory, regardless of the trajectory timescale. Rare events include ligand-receptor unbinding, protein unfolding, or any large conformational change. Rare event properties can display considerable variation, and thus should be measured as ensemble averages over many events. Therefore, for simulations to accurately characterize rare event properties in physical experiments or phenomena, they must be able to sample the property distribution resulting from multiple events. One straightforward way to sample a property distribution in

Ranjani Krishnan and Emily B. Walton contributed equally

R. Krishnan · K. J. Van Vliet
Department of Biological Engineering,
Massachusetts Institute of Technology,
Cambridge, MA 02139, USA

E. B. Walton · K. J. Van Vliet (✉)
Department of Materials Science and Engineering,
Massachusetts Institute of Technology,
Cambridge, MA 02139, USA
e-mail: krystyn@mit.edu

Published online: 06 May 2009

molecular simulations is to perform an ensemble of replicate production-run simulations, each with a different initial configuration, and observe the resulting distribution in behavior. Some researchers have begun to use multiple initial configurations for replicate simulations [4, 5]; however, two major questions have not been explicitly addressed in the literature. First, how should an ensemble of representative initial configurations be chosen? Second, how many replicate production-run simulations are necessary to characterize molecular property distributions without incurring excessive computational expense?

Selecting an ensemble of representative initial configurations generates a third, more fundamental question: can we expect there to be a discernable relationship between the minute configurational changes accessible in MD and the functional biomolecular properties of interest? The concept of a direct relationship between structure and function is essential to protein science. If configurational differences can be correlated directly with changes in properties, then a method of selecting an ensemble of initial configurations must reflect this correlation. This ensures that that the simulations utilizing these configurations will appropriately sample the property distribution of interest. However, it is unclear what kind of relationship will exist between the subtle, sub-nm configurational differences among a set of configurations selected from an ostensibly equilibrium MD trajectory and the biomolecular properties inferred from production-run MD simulation. Two recent studies highlight this question and suggest that the answer may depend on the particular biomolecular system of interest. MD simulations described by Dastidar *et al*. showed that multiple configurations of p53 peptides in complex with the protein MDM2 yield the same binding affinity [7]. In contrast, experiments reported by Coureux *et al*. showed that very small configurational differences (<0.1 nm root-mean square displacement) separate active and non-active states of a photoreceptor protein [8], indicating that photoreceptor function varies with sub-nm configurational changes.

Here, we address these three questions through a combination of clustering analysis, replicate steered molecular dynamics (SMD) simulations, and statistical analysis. We focus on a particular rare-event: unbinding of a biotin ligand from a streptavidin receptor, as studied through SMD simulations. This ligand-receptor complex of high binding affinity and long binding lifetime is widely studied through both experiments and simulations [9–15]. Unbinding occurs only once in each simulated trajectory, and our rare-event property of interest is the rupture force F_R, defined as the maximum force observed in each unbinding trajectory. We note that the loading rates accessible for large solvated proteins with current computational resources typically exceed those used in forced unbinding experiments, such that simulated F_R values cannot be extrapolated quantitatively to experimental loading rates [16]. However, computational characterization of the F_R distribution is of interest for the purposes of accurately comparing rare-event behavior under different conditions (e.g., mutated *vs*. wild-type receptors). We demonstrate a method for testing correlation between configuration and rare-event properties in any biomolecular system, and offer an approach for selecting an ensemble of initial configurations for atomistic simulation of rare-event properties in the absence of configuration-property correlation.

Methods

The biotin-streptavidin tetramer (PDB ID 1STP [17]) was simulated at 300 K in explicit water for 101 ns using the GROMACS molecular dynamics package, version 3.3 [18, 19], as described previously [20]. The portion of the equilibration trajectory with t > 15 ns was determined to have begun exploring its equilibrium phase space in at least a local minimum, as calculated by our previously described protocol [20]. Every five frames of this portion of the equilibration trajectory were used as input for clustering analysis, for a total of 489 configurations. The GROMACS tool g_cluster was used for single linkage hierarchical clustering, and clusters with at least 10 members were chosen for SMD simulations. The similarity measure for the clustering algorithm used here was the root mean square deviation (RMSD) of residues after least squares fitting between pairs of configurations. A configuration was assigned to a cluster if its RMSD with respect to any cluster member was within a cutoff, and the cluster centroid was defined as the member with the lowest average RMSD with respect to all other cluster members. The RMSD cutoff c_{RMSD} was chosen in order to produce ~10 clusters with at least 10 members each. RMSD cutoff for each clustering group: whole complex, c_{RMSD}=7.5 Å; occupied binding pocket, c_{RMSD}=3.5 Å; unoccupied binding pocket, c_{RMSD}= 2.5 Å. SMD simulations of forced biotin dissociation from the streptavidin receptor were performed as previously described [21]. The biotin was displaced *via* a Hookean spring of spring constant k = 1686 kJ mol^{-1} nm^{-2} (2.8 N/m) with a velocity $v = 0.8$ m/s

To find the mean and standard deviation σ of each distribution of rupture forces, we fit the linear portion (0-80%) of the cumulative distribution function (CDF) to a straight line. The rupture force distribution for a single ligand-receptor complex is expected to be approximately Gaussian [22], and by Taylor expansion of the Gaussian CDF, the slope of the best fit line can be approximated as

$1/(\sigma(2\pi)^{1/2})$. The mean was calculated as the force value for which the CDF equaled 50%. We were able to use the non-symmetric interval 0-80% due to the presence of a cutoff function at non-zero force caused by the added harmonic potential in *SMD* [23].

Results

Using clustering analysis to test for a configuration-property relationship

We used clustering analysis as a tool to enable testing for a correlation between initial configuration and output property F_R. Clustering analysis is a well-established tool for identifying distinct configurational groups within an MD trajectory [24–32], but to our knowledge it has not previously been applied to testing for functional differences among the configurational subsets.

All initial configurations considered in this study were selected from a single MD "equilibration trajectory" of 101 ns duration. Note that our use of the phrase "equilibration trajectory" implies only that the protein has *begun* exploring its equilibrium phase space, as determined by our previously described protocol [20] and that this equilibration trajectory is only used to generate and identify initial configurations as distinct starting points for production-run SMD simulations.

We applied a clustering algorithm to divide the frames of the equilibration trajectory into configurationally similar subsets, termed clusters. Because it is not clear which residues might capture configurational information that correlates with rupture force, we used three different groups of residues for the similarity metric in the clustering algorithm: (i) all atoms of the streptavidin tetramer with four bound biotin molecules; (ii) the streptavidin binding pocket [13, 17, 33] inclusive of the biotin ligand; and (iii) the binding pocket exclusive of the biotin ligand. The centroid (most representative configuration) of each cluster was also identified.

From the clustering results, we selected four groups of configurations to be used as input into SMD simulations: the set of cluster centroids that resulted from each similarity metric (three groups) as well as all of the configurations comprising one single cluster from each similarity metric (one group). This last criterion, including all the members of a cluster for each of the three similarity metrics, resulted in one group instead of three because it happened that this single cluster coincidentally contained the same configurations regardless of the similarity metric. The configurations within a single cluster are necessarily more similar than the set of centroids from different clusters.

All SMD simulations of forced biotin unbinding from the streptavidin receptor were conducted under otherwise identical conditions, including timescale, loading rate, spring stiffness, and spring velocity (see Methods). If the initial bound configuration were correlated with F_R of the complex, the set of all configurations from within a single cluster should result in a narrower distribution of rupture forces than the distribution observed for the set of cluster centroids (Fig. 1a). However, in our simulations, the distribution of rupture forces for a single cluster is at least as wide as that of the cluster centroids, regardless of the group used for the similarity calculation (Fig. 1b). Therefore, in this system, configurational similarity does not necessarily lead to similar rupture forces, indicating that the initial bound configuration does not correlate with or predict the force required to rupture the complex.

We also performed a set of SMD simulations using a single initial configuration but varying the randomly assigned initial velocities of the atoms (Fig. 1b). The distribution of rupture forces in this subset was approximately the same as that obtained for the subsets that differed in initial configuration, underscoring the point that initial configuration does not correlate with F_R.

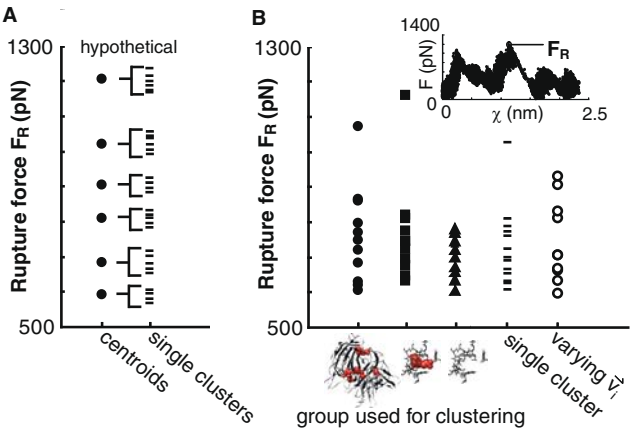

Fig. 1 (**a**) Hypothetical clustering results in the case that a correlation exists between initial bound configuration and output rupture force F_R. Filled circles denote cluster centroid rupture forces and dashes denote rupture forces for configurations within each cluster, with individual clusters indicated with brackets. (**b**) Clustering results with centroids from three different groups used for similarity calculation: Group 1: streptavidin tetramer plus four bound biotin molecules (filled circles); Group 2: streptavidin binding pocket plus one bound biotin molecule (squares); Group 3: streptavidin binding pocket (triangles). Dashes denote F_R values for configurations within a single cluster. This same set of configurations comprised one full cluster for each of the three clustering metrics used. Open circles denote F_R values for a single initial configuration with varied initial atomic velocities. Inset shows a representative force *vs*. reaction coordinate (χ) trace. Even considering three different similarity metrics, we find no evidence to support a correlation between initial configuration and rupture force

Necessary number of replicate simulations

The goal of conducting replicate atomic simulations with multiple initial configurations is to characterize a property distribution of interest. In this study, forced unbinding simulations were conducted on 53 different initial configurations (total number of configurations in the four groups selected after clustering). This is similar to the number of unbinding measurements recommended for forced unbinding experiments [34], and is assumed to be sufficient for reliable characterization of the rupture force distribution. However, it is possible that this characterization could have been done more efficiently with fewer simulations. To investigate how many replicate simulations are necessary to reliably represent the full rupture force distribution, we generated 200 subsets of N rupture force values randomly selected from the full set of 53, where N ranged from 4 to 40. We then used a two-tailed t-test to identify subsets that appeared to sample a statistically different distribution ($p<0.05$) than the full set of rupture forces and would therefore lead to incorrect characterization of the rupture force distribution. As shown in Fig. 2 for N=4, 75% of the randomly generated distributions were statistically significantly different from the full distribution, while for $N\geq30$, none of the randomly generated distributions were statistically significantly different from the full distribution.

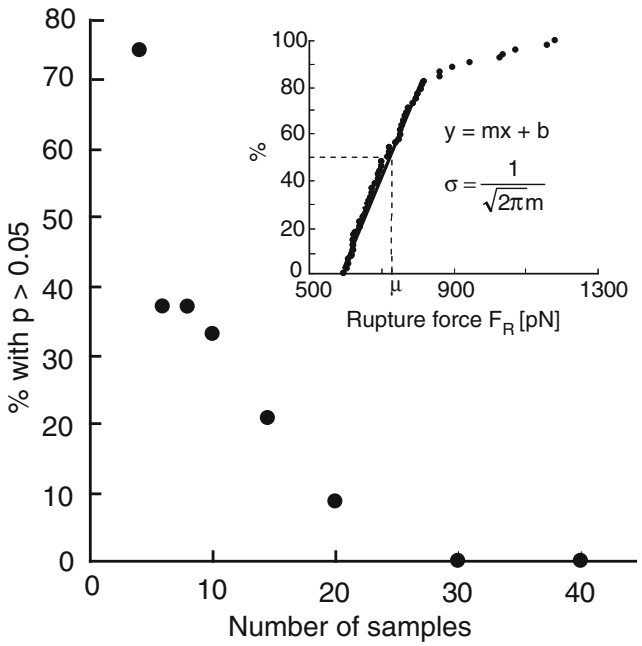

Fig. 2 The full set of 53 rupture forces was split into randomly generated subsets of size ranging from N=4 to N=40. The rupture force distribution of each subset was compared to the distribution of the full set of 53 rupture forces. The percentage of subsets that were statistically different ($p<0.05$) than the full set of rupture forces is plotted here for each N. Inset shows a representative cumulative distribution function with best fit line to the points between 0-80%. The best fit line was used to find the mean (μ) and standard deviation (σ) of the distribution as shown. Consistent with many other fields, a set of 30 samples seems to consistently reproduce features of the original property distribution

Discussion

Lack of configuration-property correlation

We found no evidence for correlation between the initial bound configuration and the measured rupture force of the biotin-streptavidin complex. However, a correlation between configuration and rare-event properties would be expected from the idealized, conceptual depiction of the energy landscapes that govern protein behavior [35]. The energy landscape along a rare-event reaction coordinate is typically depicted as smooth and one-dimensional (Fig. 3b). A small change in initial configuration would result in a small change in the energy barrier, leading to small changes in energy-dependent properties such as rupture force.

In contrast to this idealized depiction, there is evidence in the literature that the energy landscapes that govern protein behavior are highly dimensional and extremely rough on small length-scales (Fig. 3a) [15, 17, 36, 37]. The lack of correlation between initial configuration and simulated property is consistent with a rough, multi-dimensional energy landscape. Recently, Rico and Moy performed a molecular dynamics study of the biotin-streptavidin system and concluded that in a rough, one-dimensional energy landscape, complexes with similar initial configurations may have very different initial energies, leading to different energy barriers and therefore to differences in property values [15]. This is also consistent with the recent experimental study of configuration-dependent properties of photoreceptor proteins by Coureux *et al*. [8].

Here, we extend this argument to multiple dimensions, as even a rough one-dimensional reaction coordinate is not sufficient to explain the protein behavior observed in our molecular dynamics simulations. In a rough one-dimensional energy landscape, the transition state is fixed at the point of highest energy, but in multiple dimensions, the transition state is variable because there is roughness in the saddle point as well as in the energy wells. Thus, initial configurations with nearly identical initial energies can lead to trajectories with different energy barriers (Fig. 3c). Multiple dimensions also allow for variation in the direction of movement from a single configuration, meaning that identical initial configurations with different sets of initial atom velocities (i.e., multiple instances of a given configuration) may follow divergent trajectories (Fig. 3c).

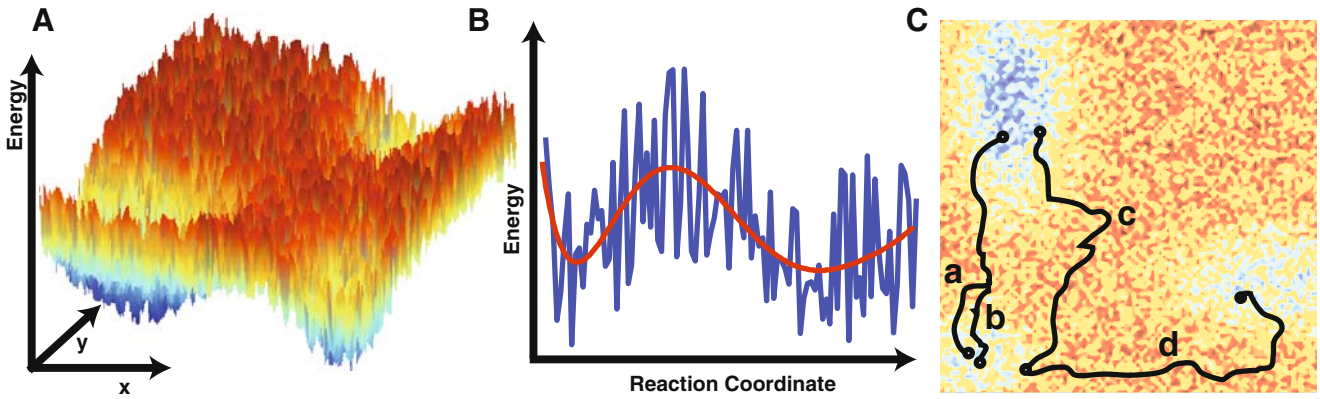

Fig. 3 Protein energy landscapes are rough and multi-dimensional (**A**, and its top-down projection, **C**), but are often represented as smooth and one-dimensional (**B**, red). In rough, one-dimensional landscapes (**B**, blue), similar configurations may have very different initial energies (pathway ***a*** and ***b*** in **C**), causing a difference in the energy barrier and therefore in values of protein properties. Multi-dimensional landscapes introduce additional mechanisms of property variation: the transition state roughness (pathways ***b*** and ***c*** in **C**) and the initial sampling direction (pathways ***c*** and ***d*** in **C**). Many protein interactions are described by rough, multidimensional energy landscapes that result in a natural variation in energy-dependent properties

Given that many protein interactions are reasonably described by a rough, multi-dimensional energy landscape [15, 27, 31, 32], we anticipate that this lack of configuration-property correlation will be observed for other rare event properties and other protein systems. However, if differences among configurations increase as phase space sampling increases (longer simulation time-scale), or if energy landscape roughness decreases (system-dependent), it is possible that correlations will exist.

Testing for a correlation between initial configuration and output property informs the selection of a representative set of configurations for production-run simulations. If configuration is not correlated with the observed property, then no particular configuration is more representative than any other; thus, *any* subset of configurations will randomly sample the underlying distribution of the observed property. For example, configurations generated from equally spaced timepoints, as used by Paci *et al*. and Curcio *et al*. [4, 5] for simulation of a hapten-antibody complex, would be a valid choice to ensure appropriate sampling.

Necessary number of replicate simulations

To efficiently characterize the underlying distribution of a biomolecular rare-event property, one should consider how many replicate samples are required. A statistical rule of thumb, used in fields as diverse as economics and epidemiology, is that approximately 30 samples are needed to adequately characterize the mean of an underlying distribution because, with $N \geq 30$, the sampling distribution approaches a normal distribution centered on the mean of the underlying distribution. This is true regardless of the shape of the underlying distribution, and has been demonstrated in a wide variety of applications [38]. Our results from randomly generated subsets of size $N=4$ to $N=40$ are consistent with this and represent the first application of this concept to molecular simulation. Although we have only explicitly demonstrated this for one particular biomolecular system, the consistency with results from other disciplines suggests that this result will

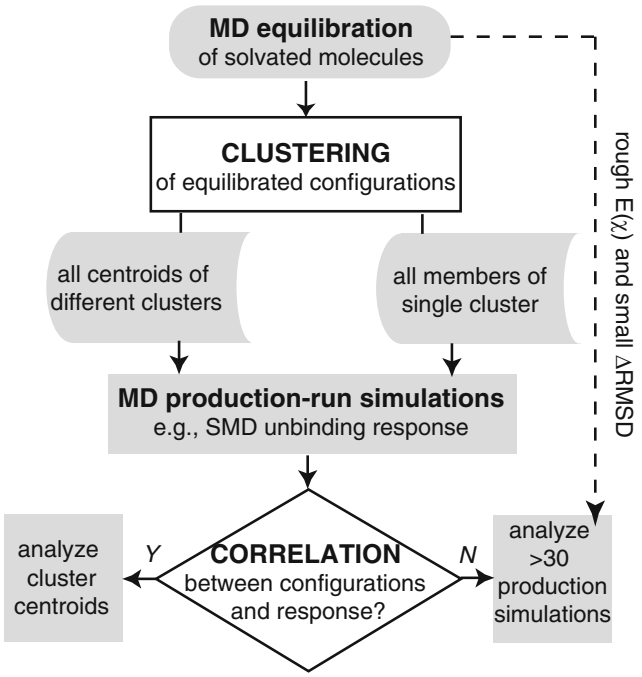

Fig. 4 A suggested method of testing for the presence of configuration-property correlation in rare event MD simulations and implications on selecting an ensemble of initial configurations for replicate production-run simulations

hold true for simulations of rare events in other biomolecular systems as well. Therefore, we predict that at least 30 replicate simulations will be necessary to ensure adequate characterization of an observed biomolecular property.

As expanding computational capabilities make it increasingly tractable to perform multiple simulations of rare-events, the atomistic simulation community should consider the procedures necessary for performing and analyzing replicate simulations. We suggest the following approach, schematically outlined in Fig. 4. First, one can use clustering analysis to test the relationship between initial configuration and the simulated rare-event property of interest, as demonstrated in this work. This includes applying a clustering algorithm to an equilibration MD trajectory, conducting production-run simulations on the set of cluster centroids and every member of a single cluster, and comparing the resulting property distributions. If there is no correlation between initial configuration and measured property (as we anticipate to be the case for protein systems with rough energy landscapes under current computational limitations), then any set of 30 configurations will randomly sample the property distribution and can be used for characterization. Note that some or all of these 30 configurations could be the same as those employed to test for a configuration-property correlation, improving computational efficiency. If 30 simulations are not computationally tractable, one can then consider means to determine appropriate confidence intervals associated with calculated property values.

However, if a configuration-property correlation exists (indicated by the observation that the property distribution exhibited by a single cluster is significantly narrower than that exhibited by the set of cluster centroids), the cluster centroids may be representative initial configurations for characterizing the property distribution of interest and must be weighted appropriately to account for cluster size. This consideration is beyond the scope of the current study.

Comprehensive atomistic simulation of rare events in biomolecular dynamics requires a new but tractable consideration of configurational replicates. The results presented here demonstrate a detailed procedure for performing replicate molecular dynamics simulations that will enable researchers to effectively characterize distributions of rare-event properties through computational simulation.

Acknowledgments The authors acknowledge the National Defense Science and Engineering Graduate Fellowship program (EBW), the National Institute Of Biomedical Imaging and Bioengineering award number T32EB006348 (RK), the National Science Foundation CAREER Award (CBET 0644846) and the Beckman Foundation Young Investigator Award (KJVV). Computational support of this project was provided by the National Science Foundation through grant IMR-0414849. We thank B. Tidor and J. M. Vinocur for insightful discussions.

References

1. McCammon JA, Gelin BR, Karplus M (1977) Dynamics of folded proteins. Nature 267:585–590
2. Maragakis P et al (2008) Microsecond molecular dynamics simulation shows effect of slow loop dynamics on backbone amide order parameters of proteins. J Phys Chem B 112:6155–6158
3. Freddolino PL et al. (2008) Ten-microsecond molecular dynamics simulation of a fast-folding WW domain. Biophys J 94:L75–77
4. Paci E et al. (2001) Forces and energetics of hapten-antibody dissociation: a biased molecular dynamics simulation study. J Mol Biol 314:589–605
5. Curcio R, Caflisch A, Paci E (2005) Change of the unbinding mechanism upon a mutation: A molecular dynamics study of an antibody-hapten complex. Protein Sci 14:2499–2514
6. Aksimentiev A, Schulten K (2005) Imaging {alpha}-hemolysin with molecular dynamics: Ionic conductance, osmotic permeability, and the electrostatic potential map. Biophys J 88:3745–3761
7. Dastidar SG, Lane DP, Verma CS (2008) Multiple peptide conformations give rise to similar binding affinities: molecular simulations of p53-MDM2. J Am Chem Soc 130:13514–13515
8. Coureux P-D et al (2008) Picometer-scale conformational heterogeneity separates functional from nonfunctional states of a photoreceptor protein. Structure 16:863–872
9. Chilkoti A, Stayton PS (1995) Molecular origins of the slow streptavidin-biotin dissociation kinetics. J Am Chem Soc 117:10622–10628
10. Chu V et al. (1998) Thermodynamic and structural consequences of flexible loop deletion by circular permutation in the streptavidin-biotin system. Protein Sci 7:848–859
11. Freitag S et al (1999) A structural snapshot of an intermediate on the streptavidin-biotin dissociation pathway. Proc Nat Acad Sci USA 96:8384–8389
12. Freitag S et al. (1997) Structural studies of the streptavidin binding loop. Protein Sci 6:1157–1166
13. Grubmüller H, Heymann B, Tavan P (1996) Ligand binding: Molecular mechanics calculation of the streptavidin-biotin rupture force. Science 271:997–999
14. Hyre DE et al. (2002) Early mechanistic events in biotin dissociation from streptavidin. Nat Struct Biol 9:582–585
15. Rico F, Moy VT (2007) Energy landscape roughness of the streptavidin-biotin interaction. J Mol Recognit 20:495–501
16. Izrailev S et al. (1997) Molecular dynamics study of unbinding of the avidin-biotin complex. Biophys J 72:1568–1581
17. Weber PC et al (1989) Structural origins of high-affinity biotin binding to streptavidin. Science 243:85–88
18. Lindahl E, Hess B, van der Spoel D (2001) GROMACS 3.0: A package for molecular simulation and trajectory analysis. J Mol Mod 7:306–317
19. Berendsen HJC, van der Spoel D, van Drunen R (1995) GROMACS: A message-passing parallel molecular dynamics implementation. Comput Phys Commun 91:43–56
20. Walton EB, Van Vliet KJ (2006) Equilibration of experimentally determined protein structures for molecular dynamics simulation. Phys Rev E 74:061901
21. Walton EB, Lee S, Van Vliet KJ (2008) Extending bell's model: How force transducer stiffness alters measured unbinding forces and kinetics of molecular complexes. Biophys J 94:2621–2630

22. Merkel R et al (1999) Energy landscapes of receptor-ligand bonds explored with dynamic force spectroscopy. Nature 397:50–53
23. Friddle RW et al. (2008) Near-equilibrium chemical force microscopy. J Phys Chem C 112:4986–4990
24. Troyer JM, Cohen FE (1995) Protein conformational landscapes: Energy minimization and clustering of a long molecular dynamics trajectory. Proteins: Struct Funct Bioinf 23:97–110
25. Li Y (2006) Bayesian model based clustering analysis: Application to a molecular dynamics trajectory of the HIV-1 integrase catalytic core. J Chem Inf Model 46:1742–1750
26. Laboulais C et al. (2002) Hamming distance geometry of a protein conformational space: application to the clustering of a 4-ns molecular dynamics trajectory of the HIV-1 integrase catalytic core. Proteins: Struct Funct Bioinf 47:169–179
27. Karpen ME, Tobias DJ, Brooks CL (1993) Statistical clustering techniques for the analysis of long molecular dynamics trajectories: analysis of 2.2-ns trajectories of YPGDV. Biochemistry 32:412–420
28. Gabarro-Arpa J, Revilla R (2000) Clustering of a molecular dynamics trajectory with a hamming distance. Comput Chem 24:693–698
29. Feher M, Schmidt JM (2003) Fuzzy clustering as a means of selecting representative conformers and molecular alignments. J Chem Inf Comput Sci 43:810–818
30. Brigo A et al. (2005) Comparison of multiple molecular dynamics trajectories calculated for the drug-resistant HIV-1 integrase T66I/M154I catalytic domain. Biophys J 88:3072–3082
31. Gordon HL, Somorjai RL (1992) Fuzzy cluster analysis of molecular dynamics trajectories. Proteins: Struct Funct Bioinf 14:249–264
32. Lei H et al. (2007) Folding free-energy landscape of villin headpiece subdomain from molecular dynamics simulations. Proc Nat Acad Sci USA 104:4925–4930
33. Lindqvist Y, Schneider G (1996) Protein-biotin interactions. Curr Opin Struct Biol 6:798–803
34. Franz CM et al (2007) Studying integrin-mediated cell adhesion at the single-molecule level using AFM force spectroscopy. Sci STKE 406:l5
35. Wales DJ (2003) Energy landscapes. Cambridge University Press, Cambridge UK, Cambridge Molecular Science
36. Honeycutt JD, Thirumalai D (1990) Metastability of the folded states of globular proteins. Proc Nat Acad Sci USA 87:3526–3529
37. Bryngelson JD et al. (1995) Funnels, pathways, and the energy landscape of protein folding: a synthesis. Proteins: Struct Funct Bioinf 21:167–95
38. Mendenhall W, Beaver RJ, Beaver BM (2006) Introduction to probability and statistics, 12th edn. Thomson Higher Education, Belmont CA