

A “black-box” re-weighting analysis can correct flawed simulation data, after the fact

F. Marty Ytreberg^{*†}, Daniel M. Zuckerman^{‡§}

[†]Department of Physics, University of Idaho, Moscow, ID 83844-0903, and [§]Department of Computational Biology, University of Pittsburgh, 3501 Fifth Avenue, Pittsburgh, PA 15260

Submitted to Proceedings of the National Academy of Sciences of the United States of America

There is a great need for improved statistical sampling in a range of physical, chemical and biological systems. Even simulations based on correct algorithms suffer from statistical error, which can be substantial or even dominant when slow processes are involved. Further, in key biomolecular applications, such as the determination of protein structures from NMR data, non-Boltzmann-distributed ensembles are generated. We therefore have developed the “black-box” strategy for re-weighting a set of configurations generated by arbitrary means to produce an ensemble distributed according to any target distribution. In contrast to previous algorithmic efforts, the black-box approach exploits the configuration-space density observed in a simulation, rather than assuming a desired distribution has been generated. Successful implementations of the strategy, which reduce both statistical error and bias, are developed for a one-dimensional system, and a 50-atom peptide, for which the correct 250-to-1 population ratio is recovered from a heavily biased ensemble.

Ensemble averages over configurations are fundamental to the analysis of finite-temperature systems of physical, chemical, and biological interest, as well as to any statistically defined system. Yet it is well appreciated that estimates of such averages based on computer simulations can suffer from both systematic and statistical error [1, 2]. We therefore ask: *Given a set of previously generated configurations of uncertain quality, what is the best way to estimate ensemble averages?* Our proposed answer, the “black-box re-weighting” (BBRW) strategy described below, appears promising in its ability to overcome both types of error in some systems.

Statistical error is a ubiquitous problem of underappreciated practical importance. Every algorithm known to the authors, including sophisticated methods [3–5], relies on repeated visits to a state (a subset of configuration space) in order to generate statistical reliability or precision in the population estimate for that state. If we define the simulation-and-system-specific correlation time t_{corr} as the time required to visit all important states at least once, then statistical precision requires a long total simulation time, $t_{\text{sim}} \gg t_{\text{corr}}$. Standard square-root-of-duration arguments [1, 2] suggest that a simulation retains a fractional imprecision of $\sqrt{t_{\text{corr}}/t_{\text{sim}}}$ (on a unit scale). Below, we show that BBRW dramatically cuts statistical error, avoiding the slow square-root behavior.

Systematic bias is typical in some systems of great practical importance—such as full-sized proteins—where, to date, it is not clear that any atomically detailed simulation has come close to reaching t_{corr} . Indeed, surprisingly lengthy simulations are required to obtain statistical ensembles for small peptides [6]. Nevertheless, biased non-Boltzmann distributed sets of atomically detailed protein structures are regularly generated, e.g., for NMR structure determination [7] and in the study of protein folding [8]. These sets can be useful for docking [9], which is employed in drug design. The proposed

BBRW strategy, in principle, can convert such sets into statistically distributed ensembles.

Background and theory

The fundamental basis of our re-weighting approach is the recognition that a simulation’s output (the configurations generated) *contains valuable information that is not conventionally used* to estimate ensemble averages. This information is the observed configuration-space density. It is critical to note that the observed density never matches the targeted distribution (e.g., proportional to a Boltzmann factor) because of (i) the omnipresent statistical error, and also possibly, (ii) errors in the algorithm or software. Our approach can treat both statistical and systematic error because it is blind to the source of error.

Theory. Standard re-weighting theory assumes that configurations in an ensemble are already Boltzmann distributed [10–14]. The novelty of our re-weighting method is that *we treat configurations as if they were generated by an unknown process*, a “black box.” Thus, any ensemble can be treated since the results of the re-weighting do not depend on the specific process that was used to generate the ensemble.

In our re-weighting approach, if the target distribution is P , then a configuration j is assigned a (relative) weight

$$W^{\text{bb}}(j) = P(j)/P^{\text{obs}}(j), \quad [1]$$

where $P^{\text{obs}}(j)$ is proportional to the observed probability density of configuration j in the ensemble generated by the simulation, discussed below. In typical cases of physical interest, P is given (exactly) by a standard Boltzmann factor. Eq. [1] is also exact, so the key to practical application of this equation is accurate estimation of P^{obs} .

The basis of Eq. [1] is the straightforward logic of re-weighting [11]. In an ideal sample, a small configuration-space volume $d\vec{x}$ about j would contain a quantity of configurations proportional to $P(j) d\vec{x}$. However, in actuality, $P^{\text{obs}}(j) d\vec{x}$ is obtained. The erroneous sampling should be corrected to cancel the error, which is what Eq. [1] does.

Ensemble averages of any function f , in our approach, are given by

$$\langle f \rangle_{\text{bb}} = \sum_{j=1}^N W^{\text{bb}}(j) f(j) / \sum_{j=1}^N W^{\text{bb}}(j), \quad [2]$$

* E-mail: ytreberg@uidaho.edu

‡ E-mail: dmz@ccb.pitt.edu

©2007 by The National Academy of Sciences of the USA

based on N configurations. The average [2] differs from what we term a “canonical average” of the same configurations, where one assumes $P^{\text{obs}}(j) \propto P(j)$; then $W^{\text{bb}}(j)$ is a constant, and averages reduce to the usual un-weighted estimate $\langle f \rangle_{\text{can}} = \sum_{j=1}^N f(j)/N$ [2]. Note that overall normalization factors (partition functions) of $P(j)$ and $P^{\text{obs}}(j)$ cancel in [2] and are not pertinent to the present report.

Several additional points should be made regarding the strategy embodied in [1] and [2]: (i) It is an *analysis* scheme, treating the same set of configuration as might be analyzed canonically. (ii) The observed probabilities $\{P^{\text{obs}}(j)\}$ will *always* differ from those intended (e.g., canonical Boltzmann factors) due to statistical error—and perhaps significantly. (iii) The weights [1] are valid regardless of the specific process used to generate configurations. (iv) The approach is particularly suited to estimating conformational free energy differences, as described below. (v) As in the single histogram approach [10], the BBRW method retains the ability to re-weight configurations into an arbitrary target ensemble P . (vi) The principal practical challenge in implementing our strategy is the estimation of P^{obs} .

Free energy differences. To provide a proof of principle for our re-weighting approach we will determine free energy differences (ΔF), between states for various test systems [15–21].

A well established canonical technique for determining ΔF is to use ordinary simulation methods (e.g., molecular dynamics) to generate an ensemble which is distributed according to the desired P . Then, the free energy difference is defined as

$$e^{-\Delta F/k_B T} \equiv Z_B/Z_A = N_B/N_A, \quad [3]$$

where N_s is the number of configurations observed in state s , and $Z_s = \int_{\vec{x} \in s} d\vec{x} \exp[-U(\vec{x})/k_B T]$ is the configurational partition function for state s with coordinates given by \vec{x} .

For our re-weighting approach, we do not assume that configurations are distributed according to P . Instead, each configuration must first be assigned a weight using Eq. [1]; then ΔF is estimated based on summing these weights in the respective states:

$$e^{-\Delta F/k_B T} \equiv \frac{Z_B}{Z_A} = \frac{\sum_{j \in B} W^{\text{bb}}(j)}{\sum_{j \in A} W^{\text{bb}}(j)}, \quad [4]$$

where $\sum_{j \in s}$ represents a sum over all weights $W^{\text{bb}}(j)$ in the state s .

One-dimensional thought experiment. The central idea behind our re-weighting approach is that one can utilize the *observed* configuration-space density to re-weight an arbitrary set of configurations into any desired ensemble. Because this idea is so fundamental to the present discussion, and apparently novel, it is useful to illustrate the idea using a one-dimensional potential.

Consider the equal-energy double-square-well system depicted in Fig. 1, for which we would like to estimate the relative populations of the two states. Assume that the barrier is so high that it cannot be crossed in a conventional simulation of affordable length, but that we can run independent simulations in each state (a common situation for biomolecular systems when multiple structures are known). The states’ widths are in the ratio 1:10.

After generating, say, equi-sized trajectories in each state, a conventional view would hold that no analysis of the relative populations is possible since the barrier has not been crossed. However, since the right well is ten times wider than

the left, the observed configuration-space density in the right well will be, on average, one-tenth that in the left. Combined with knowledge of the two ensemble sizes, this is sufficient information for calculating the state populations.

More explicitly, one can estimate the observed density by counting configurations in equi-sized histogram bins, leading to $P^{\text{obs}}(j) \propto n_b(j)$, where $n_b(j)$ is the count in the bin containing configuration j . In this example, $P(j) = \text{const}$ for all j because the energy is constant over both states. Now, by using Eq. [1] in Eq. [4], the sum over weights leads to *exact* cancellation of n_b for each bin. Thus, this purely entropic case is correctly handled, ultimately yielding simply the ratio of the numbers of (equi-sized) bins in each state—i.e., the ratio of state widths. Below, we will see that the “easier” energetic effects are also treated correctly in less trivial examples.

The key point is that *the observed density P^{obs} , combined with knowledge of the desired ensemble P provides sufficient information to deduce the state populations, regardless of the sampling bias.* Perhaps equally importantly, we note that the density can be estimated in different ways besides binning, e.g., using configuration-space distances between nearby configurations.

The reasoning used to correct systematic bias can be applied equally to statistical error. Assume now that we possess a single continuous trajectory which has crossed the barrier in Fig. 1 at least once. The population estimates will be statistically flawed due to finite sampling. Yet, our re-weighting analysis will still yield correct results since the approach does not depend on the source of the bias in the data.

Below we will show that such density information is both present and often accessible in high-dimensional systems. We emphasize that the information is always present in principle and does not depend on prior knowledge.

Results and discussion

How can the observed density be computed in a way that will be useful, even for high-dimensional systems? We will explore two distinct strategies: binning and also the use of configuration-space distances (“nearest-neighbors strategy”). For high-dimensional systems, we will demonstrate that not all coordinates need be considered in the analysis. We emphasize that other strategies for estimating density are possible [22, 23].

The test systems chosen for this study are a one-dimensional double-well potential and a 50-atom dileucine peptide (ACE-(leu)₂-NME). Our reasoning for choosing these systems are: (i) The system sizes are small enough that it is possible to compute an independent ΔF estimate via counting, i.e., using Eq. [3]. (ii) Both are non-trivial: the one-

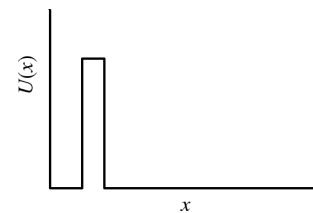


Fig. 1. One-dimensional potential for a “thought experiment.” The energy is constant and equal in both states, but the right state is 10 times as wide as the left.

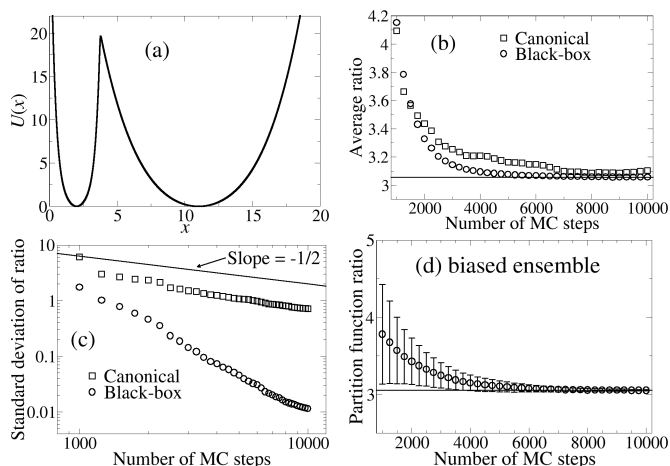


Fig. 2. Results for a one-dimensional double-well test system. (a) The potential energy plotted as a function of the coordinate. (b) The ratio of partition functions between the right and left wells. All results shown are the average of 10^3 independent computations. The canonical estimates (squares) are given by [3], the black-box estimates (circles) by [4]. An independent estimate, obtained via integration, is shown by the solid horizontal line in (b). (c) The standard deviation from the same data as that for figure (b) is plotted showing the rapid decrease in the standard deviation for our method (circles) as compared to canonical (squares). (d) Estimate of the ratio of partition functions using our re-weighting approach via Eq. [4], based on a heavily biased ensemble of configurations. Note that canonical estimation using Eq. [3] is not possible.

dimensional potential has high barrier, and dileucine has large side chains and thus 144 degrees of freedom.

Binning strategy. In the simplest use of bins, the observed density is estimated based on simple counting: omitting normalization, $P^{\text{obs}}(j) = n_b(j)$, where $n_b(j)$ is the number of configurations in the bin which includes configuration j . Indeed we have tried this scheme and it works, but not optimally.

A more effective procedure for weighting individual configurations combines simple counting with the assumption of local equilibration—within bins, configurations are assumed distributed according to P itself. This scheme corresponds to the estimate

$$P^{\text{obs}}(j) \simeq n_b(j) P(j) / \bar{P}_j. \quad [5]$$

The bin-specific normalization is the average target probability in the bin, namely, $\bar{P}_j = \sum_{i \in \text{bin}} P(i) / n_b(j)$. The approximation in [5] is that \bar{P}_j is a crude estimate of the local partition function for the bin containing configuration j . Further theoretical explanation of the approximation is provided in the Supplementary Material.

One-dimensional double-well potential. Here we consider the smooth potential of Fig. 2a. The target distribution is defined by $P(x) = \exp[-U(x)/k_B T]$, where $U(x)$ is the potential energy sketched for a $20 k_B T$ barrier height. The barrier heights for this potential, as well as basin curvatures and separations, are fully adjustable, using, $U(x) = -\ln[\exp(-0.5(x-x_a)^6 - (x-x_a)^2/2s_a^2) + \exp(-0.003(x-x_b)^4 - (x-x_b)^2/2s_b^2)]$, where $x_a = 2.0$, $s_a = 0.5$, $s_b = 1.5$. For a barrier of approximately $2 k_B T$, we use $x_b = 6.0$ and for a $20 k_B T$ barrier, we use $x_b = 11.0$. We estimated the ratio of partition functions for the two wells using data obtained from ordinary canonical simulation, i.e., Metropolis Monte Carlo (MC) [24].

Fig. 2 shows the differences between re-weighting, using a binning strategy, and canonical estimates of the ratio of partition functions between the left and right wells in (a). Results shown are the averages (b) and standard deviations (c) from 10^3 independent computations. The canonical estimates (squares) were computed via [3], and the black-box estimates (circles) via Eqs. [4] and [5] with a bin size of 0.05. All data were generated using the same MC trajectories, each initiated from the left basin of the potential in (a). An independent estimate, obtained via numerical integration, is shown by the solid horizontal line in (b). Note that in (b) the apparent accuracy of canonical estimation at short times represents only the average behavior; the statistical uncertainty is so large that any individual estimate would be completely unreliable.

Figure 2c demonstrates that the error in our re-weighting analysis decreases at a rate greater than the typical square-root behavior, as seen for the canonical error. Since the same ensembles were used for both the re-weighting and canonical estimates, this improvement may lead one to believe we are getting something for nothing. In fact, there are costs, but in essence, they have already been paid: (i) The re-weighting approach requires knowledge of the desired distribution; in our case (and most other cases of interest) this is given by the Boltzmann factor. (ii) The observed density is available information that is not generally used for analysis of simulation data, and requires only a small additional computational cost.

In Fig. 2d we show that our re-weighting approach can be used to estimate the correct populations between the left and right wells, with an example that simply cannot be treated canonically. We generated a biased ensemble, with configurations that were not distributed according to P (Boltzmann factor for U shown in (a)). Then Eqs. [5] and [4] were used with a bin size of 0.05 to estimate the ratio of partition functions. The data show the averages (circles) and standard deviations (error bars) for 10^3 independent computations. Since our re-weighting approach is independent of the method used to generate the ensemble, we are able to obtain the correct ratio, even though the ensemble used in the analysis was heavily biased. For completeness, we note that the biased ensemble used configurations generated by MC according to a flat-bottomed single-well potential spanning both states.

Dileucine peptide. In this high-dimensional molecular ex-

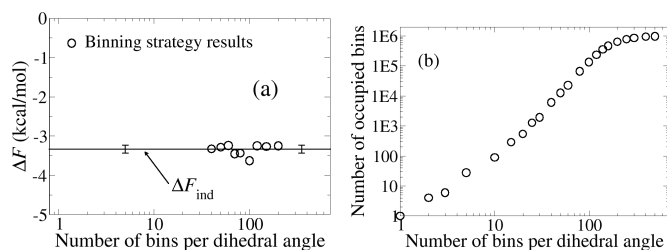


Fig. 3. (a) The free energy difference between the alpha and beta states of dileucine. Results from our re-weighting approach using a heavily biased ensembles are shown. Data were generated using Eqs. [1], [4] and [5], using only the backbone dihedrals Φ and Ψ . The independent estimate is shown by the solid horizontal line with an approximate error determined by block averaging [25]. (b) Box-counting analysis of dileucine configurations used to determine the range of bin sizes for figure (a).

ample, we again apply our approach to a biased ensemble, where canonical estimation would be meaningless. We use the method to estimate the free energy difference ΔF between conformational states of the 50-atom dileucine peptide based on independent simulations of each state. The two states considered were defined by backbone dihedral angles—“alpha”: $-140.0 < \Phi_{1,2} < -80.0$ and $-120.0 < \Psi_{1,2} < -60.0$; and, “beta”: $-125.0 < \Phi_{1,2} < -65.0$ and $120.0 < \Psi_{1,2} < 180.0$.

To test the accuracy of the results from our re-weighting approach, we also generated an independent estimate of ΔF . An ordinary, unrestrained, 1.0 μs simulation was performed and analyzed via Eq. [3], yielding $\Delta F_{\text{ind}} \approx -3.3 \pm 0.1$ kcal/mol, with statistical error estimated via block averaging [25]. The simulation data were generated using TINKER 4.2 with the OPLS-AA forcefield [26]. A 1.0 fs timestep was utilized, and the simulation was coupled to a 300.0 K bath using the Andersen thermostat [27]. Solvation was provided by the generalized Born surface area (GBSA) implicit model [28].

To generate a BBRW estimate based on the locally equilibrated scheme of Eq. [5], simulations of dileucine were performed constrained to either the alpha or beta state. We generated 500,000 configurations in each state based on a total of ≈ 104 ns. Bins used in [5] were constructed uniformly using only the four backbone dihedrals—out of 144 degrees of freedom—leading to a four-dimensional histogram.

Figure 3a shows that our re-weighting approach can be used to correctly predict the free energy difference between the alpha and beta states of dileucine. The solid horizontal line is the independent estimate ΔF_{ind} computed via Eq. [3] with an approximate error determined by block averaging [25]. The circles are results from our re-weighting method using the backbone dihedral angles Φ and Ψ . These data were computed using Eqs. [1], [4], and [5], and plotted as a function of the number of bins used per dihedral angle in the histograms. Even though an equal number of alpha and beta configurations were utilized in our analysis, the ratio of ≈ 250 is recovered. Note that our strategy, as shown here, is an example of an end-point free energy difference method, since no path connecting alpha and beta is needed (e.g., Refs. [15–21]).

Fig. 3a also demonstrates a certain robustness in the approach: there is broad range of bin sizes that can be used generate the correct free energy difference.

Two criteria were used to determine the range of bin sizes used for the re-weighting results in Fig. 3a: (i) The range of bin sizes for accurate ΔF estimation apparently corresponds to the power-law regime of a plot for estimating the box counting dimension [29]. This plot is shown in Fig. 3b where the linear regions denote the power-law behavior. (ii) The extreme bin size results are known exactly. If one bin per dihedral angle is chosen, then the estimate will always be poor and $\Delta F = 0$ (for equi-sized ensembles). At the other extreme, if a very large number of bins per dihedral angle is used, then each bin will have a count of 1, and $P^{\text{obs}} = 1$ for every structure. Since no true density information is obtained near either extreme, we excluded such estimates.

Nearest-neighbors strategy. In addition to the binning implementation above, we now estimate P^{obs} of Eq. [1] by computing distances between conformations. The motivation behind pursuing such non-binning approaches is the realization that for biomolecular systems, it may prove difficult to populate the necessary high-dimensional histogram bins.

In general, to implement the nearest-neighbors strategy, the distances between all configurations in the ensemble are computed, using any metric which preserves the Cartesian volumes intrinsic to partition functions. By definition, the dimensionality of the metric can be as large as the number of degrees of freedom needed to describe a conformation. The observed configuration-space density for structure j is then given by

$$P^{\text{obs}}(j) \propto \frac{N_{\text{dist}}(j)}{R_{\text{hyp}}(j)^d}. \quad [6]$$

Here, $R_{\text{hyp}}(j)$ is the radius of the hypersphere centered on structure j enclosing $N_{\text{dist}}(j)$ structures, and d is the dimensionality of the distance metric used. Typically, as we will demonstrate below, d can be chosen to be much smaller than the dimensionality of the system. See Supplementary Material for further discussion of the approximations entailed when d is less than the full dimensionality of the system.

We used Eq. [6] to estimate P^{obs} for a particular structure j using the following implementation: (i) Compute the distance (using any appropriate metric) between structure j and all other structures in the ensemble. (ii) Choose a value of N_{dist} . The radius of the hypersphere R_{hyp} is given by the distance to the N_{dist} nearest neighbor. For example, if $N_{\text{dist}} = 10$ is chosen, then R_{hyp} is given by the distance to the tenth nearest neighbor.

Dileucine peptide. We tested the approach embodied in Eq. [6] by estimating the free energy difference (ΔF) between the alpha and beta conformations of dileucine, using the same restrained (thus biased) ensemble that was used above. Following the previous success of using only the torsion angles of dileucine, we chose to use a dihedral angle distance between structures j and k defined by

$$D_{jk} = \sqrt{\sum_{l=1}^d (\phi_j^l - \phi_k^l)^2}, \quad [7]$$

where ϕ^l could be any dihedral angle (e.g., Φ_1 , Ψ_2 , χ_1), and d is the number of angles used in the computation, i.e., the metric dimension.

Figure 4 shows that our re-weighting approach, using the nearest-neighbors strategy, can be used to predict the free energy difference between the alpha and beta states of dileucine. The solid horizontal line is the independent estimate ΔF_{ind}

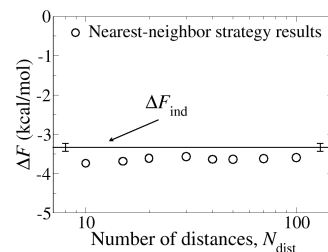


Fig. 4. The free energy difference (ΔF) between the alpha and beta states of dileucine. Results from our re-weighting approach using a heavily biased ensembles are shown. Data were generated using Eqs. [1], [4] and [6], with distances computed via Eq. [7]. The analysis used only the backbone dihedrals Φ and Ψ . Data are given as a function of the number of distances used in the analysis. The independent estimate is shown by the solid horizontal line with an approximate error determined by block averaging [25].

computed via Eq. [3] with an approximate error determined by block averaging [25]. The circles are results from our re-weighting method using the backbone dihedral angles Φ and Ψ . These data were computed using Eqs. [1], [4], and [6], with distances given by Eq. [7]. The results are plotted as a function of the number of neighbors used in the analysis procedure, N_{dist} . Even though an equal number of alpha and beta configurations were utilized in our analysis, the ratio of ≈ 250 is again recovered.

The accuracy of the nearest-neighbors strategy (or at least the way we implemented the idea) is lower than that of the binning approach; see Fig. 3a. However, the results in Fig. 4 are encouraging, and demonstrate that it is possible to estimate P^{obs} using a non-binning procedure.

Note on efficiency. It is worthwhile to briefly discuss the efficiency of the re-weighting approach using Eq. [4] compared with that of using Eq. [3]. There are, in general, two extreme cases to consider: (i) The barrier between the states of interest is very low. This implies that barrier crossing events occur frequently, and thus counting via Eq. [3] will likely be the most efficient approach. (ii) The barrier between the states is too high to cross during ordinary simulation. In this case, counting via Eq. [3] is not possible. Importantly however, the re-weighting approach of Eq. [4] can still be utilized since all that is required are independent simulations in each state.

The fairly lengthy simulations required for the BBRW method in the case of dileucine reflect the choice of state definitions: *within each state*, i.e., alpha or beta, the rotameric (side-chain) timescales are substantially longer than those of the backbone torsions.

Dimensionality. With dileucine we were able to obtain accurate results for P^{obs} , and thus ΔF , using only 4 degrees of freedom—out of a possible 144. The underlying assumption is that the degrees of freedom not included in the analysis will provide the same contribution to each weight (see Supplementary Material for details). Thus, at a minimum, the coordinates used to define the states of interest must be included in the re-weighting analysis (here, the backbone dihedrals Φ and Ψ).

A particularly promising strategy for the future is to use principal components analysis, or related methods, to suggest an optimal reduced set of coordinates [15,30,31]. These collective coordinates can then be used for binning or to determine configuration-space distances.

Conclusions

The black-box re-weighting (BBRW) strategy computes ensemble averages for any target distribution based on estimat-

ing the observed configuration-space density actually sampled in a simulation—without employing typical assumptions about sampling quality. In principle, the approach can be used to re-weight any ensemble, regardless of the process used to generate the configurations. Our data show that using BBRW can dramatically reduce both systematic and statistical error for some systems.

As a proof of principle, we used the approach to estimate free energy differences (ΔF) for non-overlapping states of one-dimensional and molecular systems, based on completely independent ensembles generated in each state. Further, when ordinary one-dimensional trajectories exhibiting transitions between states were re-analyzed with BBRW, statistical error was substantially reduced.

Practical application of BBRW hinges on estimation of the observed configuration-space density, P^{obs} , but the theoretical basis does not. We have obtained reasonable results using two different methods for estimating P^{obs} , emphasizing that *the central idea behind our re-weighting approach is independent of the specific strategy used to estimate P^{obs}* . Strategies beyond those examined in this report are available [22,23] and will be explored in future work.

There are two primary limitations of the overall approach. First, improved density estimation schemes will be needed for higher dimensional systems. Second, BBRW is an *analysis* method that re-weights configurations in *existing* ensembles; thus, it cannot predict populations for regions of configuration space not represented in the original data.

Our motivation for the black-box approach stems from the still-outstanding challenge of producing statistical ensembles for full-sized proteins. The black-box approach seems promising in this context because: (i) Non-Boltzmann-distributed sets of diverse structures can be generated using a variety of methods, including NMR structure prediction software [7], and by the addition of atomic detail to coarse models [8,32]. (ii) Local canonical sampling based on *ad hoc* starting structures is readily possible with existing software. (iii) Our own box-counting studies of proteins (unpublished) suggest that folded proteins act as systems with dramatically reduced dimensionality; see also [30,31]. We also hope that BBRW will prove useful in re-weighting implicitly solvated biomolecules into explicit solvent ensembles.

We thank Edward Lyman, Robert Swendsen, and David Zuckerman for very helpful discussions. Funding was provided by the Idaho EPSCoR, by NIH grants GM076569, CA078039, and fellowship GM073517, as well as by NSF grant MCB-0643456.

- Landau, D. P. & Binder, K. (2000) *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University, Cambridge).
- Frenkel, D. & Smit, B. (1996) *Understanding Molecular Simulation* (Academic Press, San Diego).
- Swendsen, R. H. & Wang, J.-S. (1986) *Phys. Rev. Lett.* **57**, 2607–2609.
- Berg, B. A. & Neuhaus, T. (1991) *Phys. Lett. B* **267**, 249–253.
- Lyman, E., Ytreberg, F. M., & Zuckerman, D. M. (2006) *Phys. Rev. Lett.* **96**, 028105–1–4.
- Lyman, E. & Zuckerman, D. M. (2006) *Biophys. J.* **91**, 164–172.
- Spronk, C. A. E. M., Nabuurs, S. B., Bonvin, A. M. J. J., Krieger, E., Vuister, G. W., & Vriend, G. (2003) *J. Biomolec. NMR* **25**, 225–234.
- De Mori, G. M. S., Micheletti, C., & Colombo, G. (2004) *J. Phys. Chem. B* **108**, 12267–12270.
- Knegt, R. M. A., Kuntz, I. D., & Oshiro, C. M. (1997) *J. Molec. Bio.* **266**, 424–440.
- Ferrenberg, A. M. & Swendsen, R. H. (1989) *Phys. Rev. Lett.* **63**, 1195–1198.
- Liu, J. S. (2001) *Monte Carlo Strategies in Scientific Computing* (Springer-Verlag, New York).
- Lee, H. K. & Okabe, Y. (2005) *Phys. Rev. E* **71**, 015102–1–4.
- Shehu, A., Clementi, C., & Kaviraki, L. E. (2006) *Proteins* **65**, 164–179.
- Abreu, C. R. A. & Escobedo, F. A. (2006) *J. Chem. Phys.* **124**, 054116–1–12.

15. Karplus, M. & Kushick, J. N. (1981) *Macromolecules* **14**, 325–332.
16. Xiang, Z., Soto, C. S., & Honig, B. (2002) *Proc. Natl. Acad. Sci. USA* **99**, 7432–7437.
17. Chang, C. E. & Gilson, M. K. (2004) *J. Am. Chem. Soc.* **126**, 13156–13164.
18. White, R. P. & Meirovitch, H. (2004) *Proc. Natl. Acad. Sci. USA* **101**, 9235–9240.
19. Chelvaraja, S. & Meirovitch, H. (2004) *Proc. Natl. Acad. Sci. USA* **101**, 9241–9246.
20. Ytreberg, F. M. & Zuckerman, D. M. (2006) *J. Chem. Phys.* **124**, 104015–1–9.
21. Ytreberg, F. M. & Zuckerman, D. M. (2005) *J. Phys. Chem. B* **109**, 9096–9103.
22. Silverman, B. W. (1986) *Density estimation for statistics and data analysis* (Chapman and Hall, London).
23. Scott, D. W. (1992) *Multivariate Density Estimation: Theory, Practice, and Visualization* (John Wiley, New York).
24. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953) *J. Chem. Phys.* **21**, 1087–1092.
25. Flyvbjerg, H. & Petersen, H. G. (1989) *J. Chem. Phys.* **91**, 461–466.
26. Jorgensen, W. L., Maxwell, D. S., & Tirado-Rives, J. (1996) *J. Am. Chem. Soc.* **117**, 11225–11236.
27. Andersen, H. C. (1980) *J. Chem. Phys.* **72**, 2384–2393.
28. Still, W. C., Tempczyk, A., & Hawley, R. C. (1990) *J. Am. Chem. Soc.* **112**, 6127–6129.
29. Strogatz, S. H. (1994) *Nonlinear Dynamics and Chaos: With Applications in Physics, Biology, Chemistry, and Engineering* (Addison Wesley, Reading).
30. Teodoro, M. L., Phillips, G. N., & Kavraki, L. E. (2003) *J. Comput. Bio.* **10**, 617–634.
31. Das, P., Moll, M., Stamati, H., Kavraki, L. E., & Clementi, C. (2006) *Proc. Natl. Acad. Sci. USA* **103**, 9885–9890.
32. deBakker, P. I. W., DePristo, M. A., Burke, D. F., & Blundell, T. L. (2002) *Proteins Struct. Funct. Genet.* **51**, 21–40.