



Fast and Accurate Multidimensional Free Energy Integration

Jérôme Hénin

► To cite this version:

Jérôme Hénin. Fast and Accurate Multidimensional Free Energy Integration. Journal of Chemical Theory and Computation, 2021, 17 (11), pp.6789-6798. 10.1021/acs.jctc.1c00593 . hal-03402188

HAL Id: hal-03402188

<https://hal.science/hal-03402188>

Submitted on 5 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast and accurate multidimensional free energy integration

Jérôme Hénin^{*,†}

[†]*CNRS, Université de Paris, UPR 9080, Laboratoire de Biochimie Théorique, Paris, France*

[‡]*Institut de Biologie Physico-Chimique-Fondation Edmond de Rothschild, PSL Research University, Paris, France*

E-mail: jerome.henin@ibpc.fr

Abstract

Enhanced sampling and free energy calculation algorithms of the Thermodynamic Integration family (such as the Adaptive Biasing Force method, ABF) are not based on the direct computation of a free energy surface, but rather of its gradient. Integrating the free energy surface is non-trivial in dimension higher than one. Here the author introduces a flexible, portable implementation of a Poisson equation formalism to integrate free energy surfaces from estimated gradients in dimension 2 and 3, using any combination of periodic and non-periodic (Neumann) boundary conditions. The algorithm is implemented in portable C++, and provided as a standalone tool that can be used to integrate multidimensional gradient fields estimated on a grid using any algorithm, such as Umbrella Integration as a post-treatment of Umbrella Sampling simulations. It is also included in the implementation of ABF (and its extended-system variant eABF) in the Collective Variables Module, enabling the seamless computation of multidimensional free energy surfaces within ABF and eABF simulations. A Python-based analysis toolchain is provided to easily plot and analyze multidimensional ABF

simulation results, including metrics to assess their convergence. The Poisson integration algorithm can also be used to perform Helmholtz decomposition of noisy gradients estimates on the fly, resulting in an efficient implementation of the projected ABF (pABF) method proposed by Lelièvre and co-workers. In numerical tests, pABF is found to lead to faster convergence with respect to ABF in simple cases of low intrinsic dimension, but seems detrimental to convergence in a more realistic case involving degenerate coordinates and hidden barriers, due to slower exploration. This suggests that variance reduction schemes do not always yield convergence improvements when applied to enhanced sampling methods.

1 Introduction

In molecular dynamics simulations, physical interactions intervene solely in the form of forces, which is why empirical potential energy functions are aptly nicknamed “force fields”. Energy estimation is, in that sense, a secondary concern. For the purpose of enhanced sampling, biased dynamics can be obtained through biasing forces, which may or may not derive from a biasing potential. Whereas Adaptive Biasing Potential methods^{1–5} involve the estimation of a free-energy surface (FES), Adaptive Biasing Force (ABF) methods⁶ are built around the estimation of the free energy gradient, in the spirit of Thermodynamic Integration (TI).⁷ A key difference between those two perspectives is that the free energy gradient is a local, absolute quantity, whereas free energy measurements are relative and hence non-local. Due to the local character of ABF and the absence of “fill rate” measuring the steady-state work exerted by the bias, the simulation time necessary to cross a barrier is independent from its height.

Usually, however, the quantity of interest is not the free energy gradient, but the free energy surface itself. In addition, the mean force used to estimate this quantity in multidimensional ABF^{8,9} is not exactly the gradient of any scalar potential, although it converges towards such a gradient. This makes estimating the free energy a non-trivial problem.

In this contribution I present the implementation of a solution to this problem based on original work by Lelièvre and coworkers.^{10,11} The implementation is part of the Collective Variables Module,¹² The formalism expresses the free energy surface as solution to a Poisson problem, making it amenable to robust numerical schemes. I also present a “real-world” implementation of projected ABF (pABF) proposed by the same authors, and use it to assess its practical merits when sampling biomolecular systems. The results are found to vary greatly between a low-dimension model and a more realistic system exhibiting many coupled slow degrees of freedom, in a similar fashion to large biomolecular assemblies.

2 Theory and methods

Consider a classical-mechanical system with Cartesian coordinates q , subject to the potential energy function $V(q)$, and evolving under Langevin dynamics or another type of thermostatted classical dynamics that samples from the canonical ensemble at temperature T , or inverse temperature $\beta \equiv 1/(k_B T)$. That is, the limiting distribution of the dynamics is the Boltzmann distribution $\rho(q) \propto \exp(-\beta V(q))$. In the following we will focus on configurational statistics, rather than dynamical properties. Given an arbitrary set of collective variables $z = \xi(q)$, the free energy surface associated to z is defined, up to an additive constant, by:

$$A(z) \equiv -kT \ln \int e^{-\beta V(q)} \delta(\xi(q) - z) dq \quad (1)$$

That is, $A(z) = -kT \ln \rho(z)$, where $\rho(z)$ is the marginal probability density of variable $z = \xi(q)$.

2.1 ABF yields estimates of the free energy gradient

ABF was originally described and implemented for a unidimensional collective variable.^{6,13} ABF in dimension greater than one^{8,9} is built around an estimator of the free energy gradient along the chosen collective variables. The implementation of ABF within the Colvars

Module¹² implements a formalism derived by Ciccotti et al.,¹⁴ based on projection of atomic forces in Cartesian coordinates onto the collective variables. The gradient of A may be estimated as:

$$G(z) = \left\langle \nabla_q V(q) \cdot \mathbf{v}(q) - \frac{1}{\beta} \nabla_q \cdot \mathbf{v}(q) \right\rangle_z \quad (2)$$

where $\mathbf{v}(q)$ is an arbitrary vector field satisfying certain orthonormality conditions with the collective variable gradient $\nabla_q \xi$, and the brackets indicate a canonical average conditioned by $\xi(q) = z$.^{9,14} Empirically, Thermodynamic Integration has been found to be more efficient than other free energy estimators, regardless of the method used to perform sampling;^{12,15,16} yet no theoretical justification has been put forward.

If force projection is not practical, the extended-system ABF (eABF) approach can be used, which estimates the gradient of a biased free energy surface A^k .^{17,18}

$$A^k(\lambda) \equiv -kT \ln \int e^{-\beta V(q)} e^{-\frac{\beta k}{2} (\xi(q) - \lambda)^2} dq \quad (3)$$

Note that Equation 3 is related to Equation 1 simply by substitution of the Dirac distribution with a Gaussian kernel, of variance inversely proportional to the coupling constant k . eABF can be complemented with unbiased estimators of the free energy gradient $\nabla A(z)$, such as Corrected z-averaged restraint (CZAR)¹⁸ or Umbrella Integration.^{19–21}

The initial implementation of ABF in the Colvars Module⁹ only yielded free energy gradient estimates: multidimensional free energy surfaces were integrated using a separate tool (`abf_integrate`), which performs discrete Markov-Chain Monte-Carlo (MCMC) sampling with a metadynamics-like history-dependent bias. That method is algorithmically simple, general, and importantly, implemented in arbitrary dimension. However, its convergence properties are unfavorable and difficult to tune. The estimate tends to fluctuate around the optimal value of the free energy, which generally results in noisy estimates, and the bias increments can be made to decay to zero over time, but this introduces an additional parameter. The present, fully deterministic approach solves these issues, improving both the

performance and reliability of the integration, and simplifying its use by eliminating most tunable parameters.

2.2 Free energy surface as “integral” of a non-conservative vector fields

Considering a **multidimensional** collective variable $z = \xi(q)$, we note $G_t(z)$ the estimate at time t of the free energy gradient. Obtaining the corresponding FES $A_t(z)$ is not as simple a task as integrating a scalar function. The main reason for that is that for finite t , G_t is not a gradient. Due to statistical noise, it is not conservative, meaning that its integral over closed curves is not zero, and as a result it does not derive from any potential (scalar field). Therefore the intuitive differential equation

$$\nabla A_t = G_t \tag{4}$$

admits no solution for A_t . Therefore, the question “*What FES is G_t the gradient of?*” is not a well-posed problem.

2.2.1 Poisson integration

Lelièvre and coworkers have proposed^{10,11} to express the free energy as the solution to a well-posed Poisson problem, by noting that the divergence of G_t is an estimator of the Laplacian of the free energy:¹¹

$$\nabla^2 A_t = \nabla \cdot G_t \tag{5}$$

The present algorithm consists in a finite-difference solution of Equation 5 on a regular grid. The solution A_t minimizes the L^2 distance $\|\nabla A_t - G_t\|_2$,¹¹ that is, it answers the question “*What FES has a gradient that is as close as possible to G_t ?*” (in the least-squares sense). Equation 5 only defines A_t up to a linear function, so it needs to be complemented

by appropriate boundary conditions, as discussed below.

Other methods to reliably reconstruct a free energy surface from gradient estimates differ by the basis functions used and the optimization method. In particular, the minimization of $\|\nabla A_t - G_t\|_2$ can be expressed directly as a set of linear equations involving the parameters of the free energy surface.²² This is used in the On-The-Fly Parameterization (OTFP) method,²³ which has been implemented up to dimension 2 on a basis of *chapeau* functions.²⁴ Expressing the minimization problem as stated in OTFP in our numerical context leads to the same expressions as when starting from the Poisson problem. Maragliano and VandenEijnden have proposed to use a variational optimization on a basis of radial functions.²⁵ A Bayesian strategy that integrates any knowledge about the error in gradient data is to apply Gaussian Process Regression,²⁶ which has been coupled with ABF.¹⁶ The present formulation of finite-difference Poisson integration is ideally suited to gradient data from ABF, which is estimated on a regularly-spaced dense grid.

2.2.2 Finite-difference Poisson integration algorithm

Integrating Equation 5 requires two numerical steps: (1) compute the discretized divergence of the current gradient estimate G_t , completed with relevant boundary conditions, and (2) solve the resulting discrete Poisson problem using a Conjugate Gradients (CG) linear solver.^{27,28} The conjugate gradients method is particularly efficient for sparse, symmetric, positive-definite linear systems, which is the case of the discrete Poisson equation. In CG, the target vector is optimized in steps along a succession of mutually orthogonal search directions. A complete yet accessible (even entertaining) introduction to CG can be found online, in a little gem of a paper by Jonathan Shewchuck.²⁷ The version of CG used here does not involve preconditioning. Tests using the universal but simple-minded Jacobi preconditioning yielded no improvements. While preconditioning is generally advised to improve the convergence of CG, it is still a subject of active research in the context of the Poisson equation with Neumann BCs,²⁹ let alone mixed boundary conditions.

In the Colvars implementation of ABF and eABF-CZAR, the conditional average used to estimate the free energy gradient G_t is collected within discrete bins, therefore the estimate is really the average gradient within the bin boundaries. In 1D, integration is performed using a simple middle Riemann sum, yielding free energy values at the bin edges. Because the gradient is a bin average accounting for all gradient values within the bin interval (Figure 1), this is the most accurate scheme. This is in contrast to purely local estimates of the gradient that could be obtained from constrained simulations; in that case, a higher-order integration scheme would be warranted to account for intermediate values of the gradient.

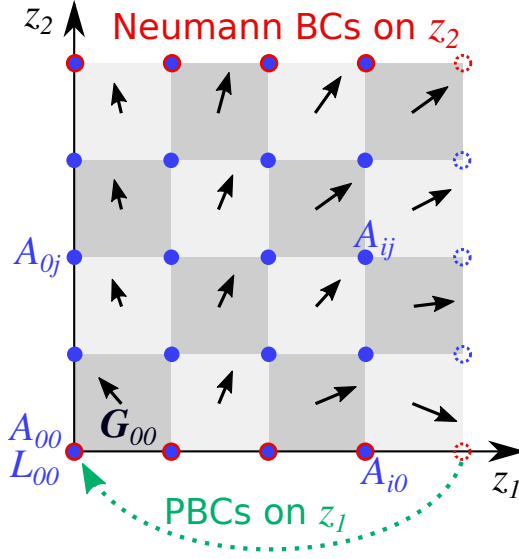


Figure 1: Two types of grids used to represent the free energy surface and its derivatives. The inner grid (shaded squares) is used for the discretized free energy gradient \mathbf{G}_{ij} (arrows), averaged over bins. The outer grid (blue dots) is used for both the discretized free energy A_{ij} and Laplacian L_{ij} . This example shows mixed boundary conditions: periodic along z_1 , and non-periodic (Neumann) along z_2 .

In accordance with that strategy, in the present scheme the free energy and its Laplacian are described on the same grid (blue dots in Figure 1), whereas the gradient grid is shifted by half a grid interval in each dimension (shaded squares and arrows in Figure 1):

$$z_i^k = z_i^0 + k\delta z_i \quad (6)$$

$$z_i^{\text{grad},k} = z_i^0 + \left(k + \frac{1}{2}\right) \delta z_i \quad (7)$$

As a result, the FES grid is larger than the gradient grid by one bin, except in PBC, where the FES grid wraps around and the last point is represented by the first point. The Laplacian grid always has the same size as the FES grid. This way, numerical differentiation can be performed as a centered two-point difference (or 2^d -point in dimension d), and the discretized Laplacian takes a classic form, using a five-point stencil in 2D and a seven-point stencil in 3D.

$$L_{ij} = \frac{A_{i-1,j} + A_{i+1,j} - 2A_{i,j}}{\delta z_1^2} + \frac{A_{i,j-1} + A_{i,j+1} - 2A_{i,j}}{\delta z_2^2} \quad (8)$$

Periodic boundaries are specified by wrapping the discrete Laplacian stencil around the boundaries. In non-periodic cases, Neumann boundary conditions are applied, that is, the value of the gradient in the normal direction to the boundary is imposed. This choice is natural given the data, which specify the gradient values everywhere (including the normal gradient at the boundaries). In contrast, classic boundary conditions such as Dirichlet cannot be applied here because there is no way to specify a priori any boundary values of the free energy itself. In that non-periodic case, the boundary element of the “Laplacian” matrix is written, taking as an example the $z_2 = z_2^0$ boundary:

$$L_{i0} = \frac{1}{2} \frac{A_{i-1,0} + A_{i+1,0} - 2A_{i,0}}{\delta z_1^2} + \frac{1}{\delta z_2} \frac{A_{i,1} - A_{i,0}}{\delta z_2} \quad (9)$$

Note that the z_1 term of the Laplacian is halved on the z_2 edge, and the z_2 term is replaced by the finite-difference normal gradient, *further divided by* δz_2 . These two modifications combined preserve the symmetry of the discrete Laplacian.^{30,31} Recall that the free energy and laplacian grids are each treated as a vector, so using the present notations, the discrete Laplacian operator is symmetric iff the coefficient of A_{ij} in L_{kl} is the same as that of A_{kl} in L_{ij} for all (i, j) and (k, l) .

Semi-periodic conditions may occur if not all colvars are periodic. They are treated consistently with the two cases above, again, keeping the discrete Laplacian matrix symmetric.

Equation 5 is solved using a conjugate gradients algorithm, taking into account the symmetry of the discrete Laplacian. This only requires the Laplacian to be expressed as a left-multiply function, which applies the Laplacian operator to a given vector. As a result there is no need to explicitly store the discrete Laplacian matrix, which is very sparse and of high dimension especially in the 3D case (dimension $(l \times m \times n)^2$, where l, m, n are the grid sizes). Therefore an efficient implementation can be written without dependence on sparse matrix storage and manipulation algorithms.

2.2.3 Practical use of Poisson integration in the Colvars Module

Within the ABF⁹ and eABF-CZAR¹⁸ algorithms in Colvars, integration of 2D and 3D FES is enabled by default. This has been included in the Colvars Module since version 2017-12-14, in NAMD 2.13 and LAMMPS version stable_16Mar2018. It is controlled by the `integrateTol` and `integrateMaxIterations` keywords, which set the tolerance and maximum number of iterations of the conjugate gradients solver.

Poisson integration routines are implemented within the generic `colvargrid` class of the Colvars Module, used to handle quantities discretized on regular grids. Thus they can be re-used by any algorithm that follows a Thermodynamic Integration-like formalism to estimate a multidimensional free energy gradient, such as Umbrella Integration¹⁹ or the combination of metadynamics and Thermodynamic Integration.¹² In particular, I provide a lightweight, standalone command-line tool `poisson_integrator` to parse gradient grid files and perform the integration. It can be used to integrate any gridded gradient estimate obtained using unrelated algorithms.

In particular, the ABF/Colvars implementation can be used to merge data from independent simulations covering different regions of colvar space (*stratification*), which may be overlapping or simply adjacent. This is done by providing data from several previous ABF simulations as input (using the `inputPrefix` keyword), and running a single timestep of simulation to trigger an ABF update. In that case, gradient estimates from all regions are

combined by ABF, then integrated using the present algorithm, which results in a single integrated FES covering the complete domain of colvars.

2.3 Projected ABF for variance reduction in real-world simulations

2.3.1 Helmholtz projection of the numerical mean force

Alrachid and Lelièvre proposed¹¹ to use the numerical gradient of the estimated FES A_t , ∇A_t , as the basis for the biasing force, thus performing a Helmholtz projection of G_t onto the space of irrotational vector fields, and eliminating its non-conservative components. Thus the “mean force” estimate of the gradient is replaced with the closest quantity (in the least squares sense, see section 2.2.1) that behaves like a true gradient, i.e. a conservative vector field. The spurious degrees of freedom present in the mean force are eliminated when computing the discrete Laplacian as part of the Poisson problem.

This pABF approach is an instance of ABF with kernel-based force reconstruction, like ABF(GPR),¹⁶ where the free energy surface is reconstructed via Gaussian Process Regression.²⁶ Using a proof-of-concept implementation, Alrachid and Lelièvre performed numerical experiments of pABF on a toy model: repulsive particles in dimension two, with a double-well potential on a single degree of freedom. This demonstrated reduced variance of the biasing force in pABF, resulting in faster convergence than standard ABF dynamics. However, the implementation did not support molecular simulations, and consequently, did not allow for numerical tests of pABF on more realistic and challenging systems.

2.3.2 Optimized on-the-fly integration for pABF

The pABF implementation of Colvars relies on Poisson integration as described above. The integrated FES is updated with at finite time intervals, and the biasing force is calculated as a finite-difference gradient of the latest integrated FES. Since the gradient and FES grids are shifted by half a bin width, the numerical differentiation is a 2^d -point centered finite difference.

The Poisson integration algorithm as described above is well-suited for a single, isolated free energy estimation. However, for pABF in particular, the FES estimate must be updated many times at short time intervals. Furthermore, the values of G_t vary more and more slowly as the algorithm approaches convergence; then, most of the work in successive updates becomes redundant. In pABF/Colvars, this redundant work is avoided in two ways: the divergence $\nabla \cdot G_t$ is updated locally where new samples have been collected. Then, at pABF update times, the conjugate gradients optimizer is restarted from the previous value of A_t , leading to rapid convergence. Finally, the optimization is performed alternately at two different precisions: at high precision when (re-)initializing the algorithm or outputting the FES, and at a lower precision upon frequent updates for pABF.

Accordingly, pABF in Colvars is controlled by the following user parameters:

`pABFIntegrateFrequency`, `pABFIntegrateTol`, and `pABFIntegrateMaxIterations`.

3 Numerical methods

Numerical tests of the Poisson integration routines shown in Figures 2, 3 are based on the following model free energy surfaces:

$$A^{2d} = \frac{1}{2}(\sin(x) \cos(2y) + 1) \tag{10}$$

$$A^{3d} = \frac{1}{2}(\sin(x) \cos(2y) + 1) + \cos(z) + 1 \tag{11}$$

The gradients of those surfaces were calculated on regular grids by Python scripts, and fed to the Poisson integration algorithm. The resulting integrated “free energy” surfaces were then compared to their exact values.

Convergence of the integration was monitored using the relative error between the computed discrete Laplacian and its target value $\|\nabla^2 A - \nabla \cdot G\|_2 / \|\nabla \cdot G\|_2$ (Equation 5), and the root-mean-square deviation (RMSD) between the integrated FES and the exact FES (Equations 10 and 11). The relative error is always available and is the stopping criterion of

the conjugate gradients solver, whereas the RMSD to the exact free energy RMSD is only available for synthetic gradient data.

The test systems considered for ABF and pABF simulations were two small peptides in vacuum that were used in many studies before: the dipeptide mimic N-acetyl-N-methyl-L-alanylamide (NANMA), and the helical peptide deca-alanine, which exhibits a complex conformational landscape.⁹

Simulations were run using NAMD 2.14³² coupled with the current version of the Collective Variables Module.¹² In vacuo systems were modeled by the CHARMM22 force field, and simulated with a base timestep of 0.5 fs and outer timesteps of 2 and 4 femtoseconds for non-bonded and long-range forces respectively. Langevin dynamics at 300 K was run with a damping coefficient of 5 ps⁻¹.

In the NANMA example, ABF dynamics was applied to the Ramachandran angles φ and ψ , with a bin width of 5 degrees. The `fullSamples` parameter was set to 20 time steps.

To compare standard ABF followed by Poisson integration with a FES obtained directly, well-tempered metadynamics simulation³³ was run on the same coordinates using the Colvars Module. The metadynamics hill height was 0.1 kcal/mol, the hill width was 9 degrees, and the hills were added every 50 timesteps, with a bias temperature of 6000 K. The biasing potential was projected onto a grid with 5-degree spacing to allow for direct comparison with ABF/Poisson. 20 replicas were run for 100 ns each, and the resulting FES were averaged to give the comparison FES. Similarly, the free energy gradient from 20 independent ABF simulations were averaged, then integrated to give the ABF FES shown in Figure 2 (right panel).

ABF simulations on deca-alanine use dihedral PCA³⁴ (dPCA) modes as collective variables. dPCA was performed using the software Carma³⁵ on a dataset of 7 structures, representing the metastable conformations identified in reference 9. `fullSamples` was set to 200. In all pABF simulations, the integrated free energy was updated every 200 steps, with a tolerance of 10^{-4} and a maximum of 100 conjugate gradient steps. In simulations of larger,

solvated systems, the integration cost would be comparatively negligible.

Convergence of free energy gradient estimators was assessed as follows. For each set of parameters, ten replica simulations were started using pseudo-random initial velocities and different seeds for the pseudo-random generator providing Langevin stochastic forces. Values of the estimators to be tested were saved regularly using the `historyFreq` parameter of ABF/Colvars.

The average of final estimates of the free energy gradient from all replicas was used as reference for measuring convergence. The gradient and free energy estimates of each replica were compared to this reference using an RMSD distance. The average and standard deviations of these distances over the ten replicas were used to assess the accuracy and rate of convergence of estimators, and are plotted in Figures 5 and 6.

4 Results

4.1 Robust and accurate free energy landscapes in dimension 2 and 3

The left panel of Figure 2 illustrates the accuracy of Poisson integration on an analytical 2d potential, in a non-trivial case of mixed boundary conditions: periodic boundary conditions on the first coordinate, and Neumann boundary conditions on the second. The 3d FES defined by Eq 11, also in semi-periodic boundary conditions, is represented in Figure 3 together with the convergence graph. As the conjugate gradients solver converges, the error between the integrated FES and the exact expression reaches a plateau at 5×10^{-6} , which is the residual due to the discretization error.

The right panel of Figure 2 illustrates the accuracy of the integrated FES on a real rather than synthetic data: the 2d FES of NANMA, with a reference surface computed using well-tempered metadynamics simulations. The RMSD between the two surfaces is only 0.1 kcal/mol for free energy values ranging from 0 to 22 kcal/mol, and can be ascribed

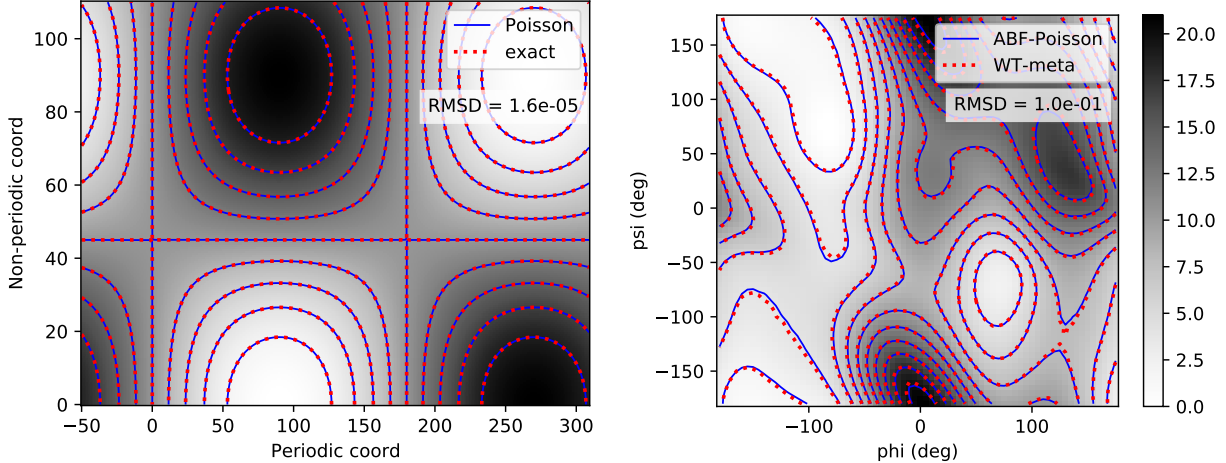


Figure 2: Left: Integration of a model 2d free energy surface (Eq. 10) from its discretized gradient in semi-periodic boundary conditions, using a discrete Poisson equation. Level lines of the exact free energy surface (blue, solid) are superimposed with those of the result of Poisson integration (red, dotted). Free energy values are given by the grayscale background. The first coordinate is periodic, whereas the second is not. The free energy scale is 0 to 1. Right: Comparison of the level lines for the free energy surface of NANMA, from ABF simulations with Poisson integration and from independent well-tempered metadynamics simulations. The free energy is given by the color scale (unit: kcal/mol).

to residual statistical variance in both free energy estimates, as well as discretization error in ABF and the finite width of Gaussian kernels in metadynamics.

Free energy gradients in dimension higher than 3, which are rarely used, can still be integrated using the existing MCMC integration tool in arbitrary dimension.⁹

4.2 Noise reduction by gradient projection in pABF

Figure 4 illustrates the Helmholtz decomposition of noisy gradient data from an ABF simulation of the NANMA system. The top left panel shows the estimate of the gradient G_t from a short, unconverged ABF simulation. The middle left panel is the conservative component of G_t . It has zero curl, and its divergence is the Laplacian of the free energy. This vector field preserves the features of G_t corresponding to variations of the underlying free energy, but is smoother. The difference $G_t - \nabla A_t$ (bottom left) is a purely non-conservative field with zero divergence and nonzero curl. It appears very noisy, and the only discernible features are

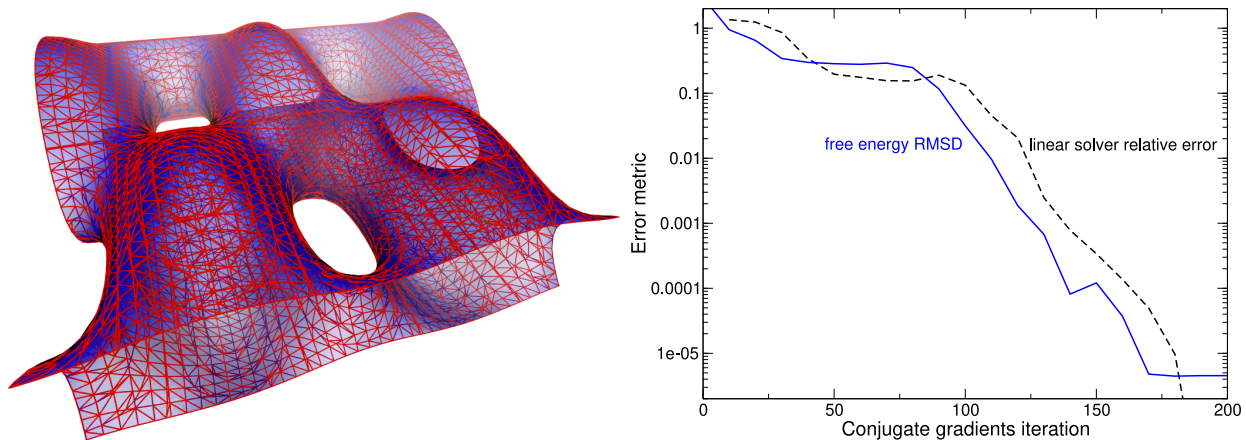


Figure 3: Left: Poisson integration of 3D gradients in semi-periodic boundary conditions. Isosurface of a model 3d free energy landscape (Eq. 11), solid blue) and the same isosurface for the potential integrated from its discretized gradient (red mesh). The two horizontal coordinates are periodic, while the vertical one follows Neumann boundary conditions. Graphics rendered using VMD.³⁶ Right: Convergence of Poisson integration by conjugate gradients. The relative error of the linear solver is represented by the dashed black line, and the free energy RMSD from the reference is shown as a solid blue line.

loops, which attest to its non-conservative character. Taken together, these illustrate that the conservative component ∇A_t is a reduced-variance estimate of the free energy gradient, compared to the ABF estimate G_t . As the ABF simulation progresses, G_t converges towards a conservative vector field, and the ABF and pABF forces tend to become identical. This is visible in the right column of Figure 4, where the same quantities are plotted for a later point in the ABF simulation, much closer to convergence. The pABF force at 3 ns (middle left), in the visited regions, is closer than the mean force (top left) to the converged mean force (top right). Thus, within sampled regions, the conservative biasing force of pABF converges faster than the mean force of standard ABF.

4.3 Rate of exploration and convergence of pABF

Convergence of free energy calculations is usually monitored using relative metrics, measuring convergence towards a known target estimate of the free energy or its gradient, or in the absence of such a reference, towards the final estimate of the present calculation. Convergence in terms of gradient is a more direct measure of the output of ABF, however it is sensitive

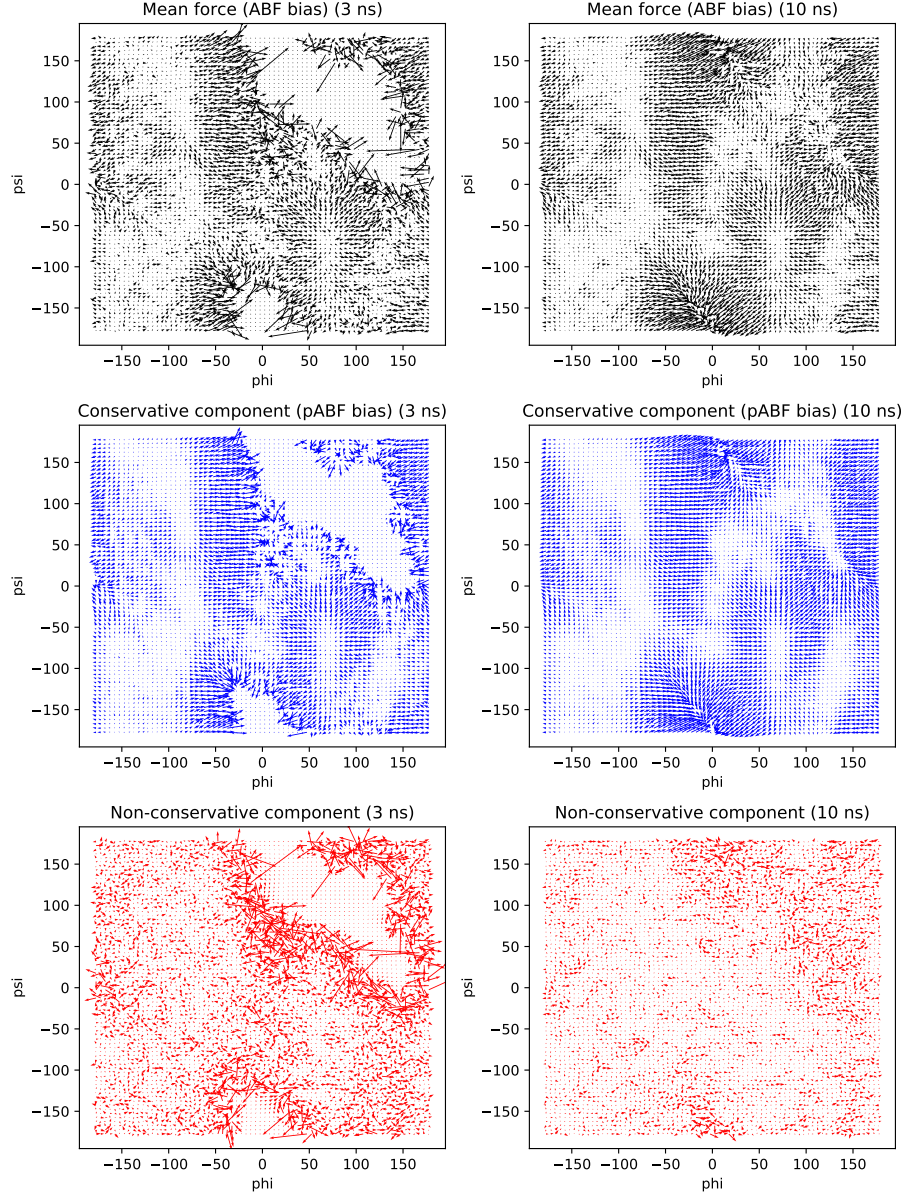


Figure 4: Helmholtz decomposition of a noisy discrete estimate of the free energy gradient for the Ramachandran angles of NANMA, in an ABF simulation after 3 ns (left column) and 10 ns (right column) of sampling. Top (black): mean force estimate G_t of the free energy gradient from ABF simulations. Middle (blue): finite-difference gradient ∇A_t of the free energy surface obtained from G_t by Poisson integration. Bottom (red): residual $G_t - \nabla A_t$ representing the non-conservative component of the gradient estimate. Arrow scale is twice that of the upper panels to enhance the visibility of small vectors.

to noise in the gradient estimates: therefore methods that yield less accurate but smoother gradient values fare better by that metric, even if the corresponding free energy is more biased. Convergence in free energy, conversely, requires an arbitrary anchoring of the free

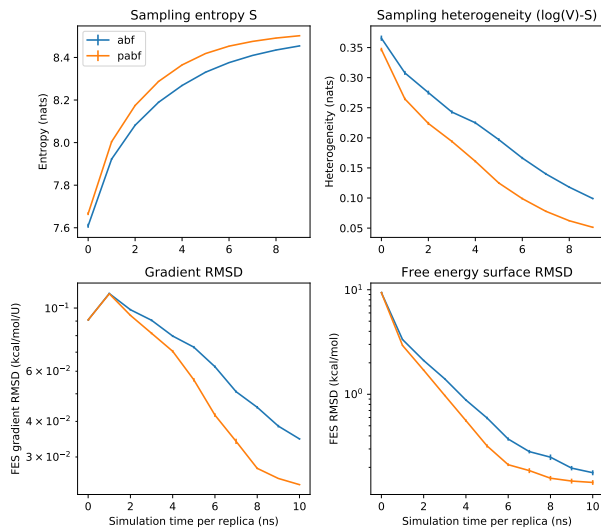


Figure 5: Exploration and convergence metrics for the 2d NANMA example (no hidden barriers), with and without pABF.

energy surfaces. Here the minimum value is set to zero before computing the RMSD.

To complement these relative metrics that depend on a known reference, I use absolute metrics of sampling in colvar space. A pure exploration metric is the natural logarithm of the visited volume, measured as the number of bins with nonzero samples: $\log(V) \equiv \log(n_{\text{visited}})$, ignoring the additive constant corresponding to the log of the bin volume.

Since ABF converges towards uniform sampling at long times, the heterogeneity of sampling can be used as an absolute convergence metric. First the sampling entropy S is calculated as:

$$S = - \sum_{i=1}^{n_{\text{visited}}} \frac{n_i}{N} \log \left(\frac{n_i}{N} \right) \quad (12)$$

where n_i is the number of samples collected in bin i and $N = \sum_i n_i$, and the sum is calculated over visited bins only. For a uniform distribution, $S = \log(n_{\text{visited}})$. Therefore, $h \equiv \log(n_{\text{visited}}) - S \geq 0$, a value of 0 corresponding to uniform sampling. I use the quantity h as a measure of the heterogeneity of sampling. These metrics are implemented in the Python module provided as Supplementary Material.

These convergence and exploration metrics for ABF on NANMA are shown in Figure 5. ABF and pABF exhibit qualitatively similar exploration and convergence profiles, but pABF

explores faster and more uniformly, and converges faster to the reference free energy profile and gradient. Therefore, the reduced-variance biasing force of pABF is beneficial to convergence on the NANMA example.

Since NANMA is intrinsically low-dimension, it is better viewed as a toy model than a relevant benchmark for sampling algorithms. Previous work has shown the multimodal character of deca-alanine conformations in vacuum,⁹ making it a minimal yet demanding system for conformational sampling. The present ABF simulations use the first 3 eigenmodes from dihedral principal component analysis (dPCA).³⁴

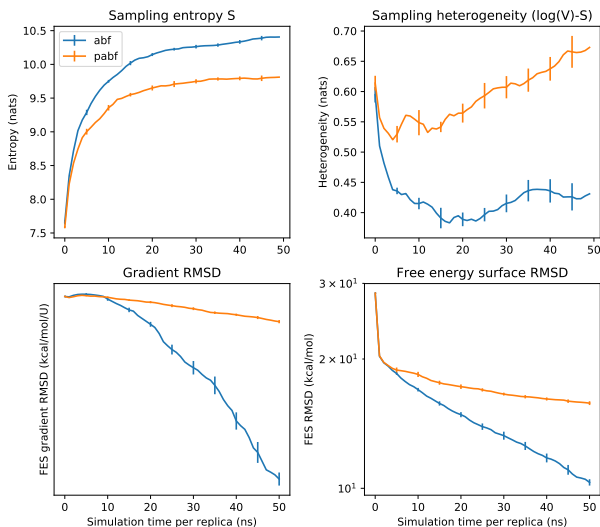


Figure 6: Exploration and convergence metrics for the 3D deca-alanine example (hidden barriers), with and without pABF.

As in the previous example, ABF and pABF converge to the same 3d free energy landscape for deca-alanine (Figure 6), with a global RMSD smaller than 2.1×10^{-1} kcal/mol from the same reference, indicating a clear consensus. In contrast with the convergence results of Figures 4 and 5, however, pABF does not exhibit improved convergence over standard ABF, by any metric. The salient feature of pABF dynamics on this system is that its exploration is considerably slower than standard ABF. This is in sharp contrast with both the present results on the simpler NANMA example, and the results of Alrachid and Lelièvre on a model system. Those two systems share the property of possessing an intrinsically low-dimension

slow space (2d), without significant “hidden barriers” along other degrees of freedom – the “orthogonal space”.^{20,37}

An unexpected result is that reducing the variance of the biasing force does not always lead to faster convergence of ABF. The presence of noise in the biasing force in the transient regime of ABF may enhance diffusion, leading to faster exploration. It might explain the efficacy of “classic” metadynamics² for exploration, with its ever-varying biasing potential that promotes constant motion in collective variable space. “Well-tempered” metadynamics³³ is a compromise between this initial behavior and long-term convergence of the free energy surface, which can also be seen as switching from exploration to exploitation. This is also exploited in the *meta-eABF* and *WTM-eABF* methods, combinations of metadynamics and ABF.³⁸ Another possibility is that classic ABF is more responsive than pABF to local fluctuations of the mean force corresponding to fluctuations in the orthogonal space,³⁹ leading to more efficient crossing of hidden barriers during the exploration phase.

pABF has common properties with ABF(GPR):¹⁶ in ABF(GPR), the FES is reconstructed by Gaussian Process Regression, then differentiated to obtain the biasing force, which is constrained to be a true gradient, as the biasing force of pABF. ABF(GPR) was shown to explore the 2d FES of NANMA significantly faster than standard ABF, and approximately as fast as a constant bias exactly compensating the underlying FES. However, the same study showed that the metadynamics leads to even faster exploration, presumably due to its self-avoiding behavior, regardless on the presence of free energy barriers, or lack thereof.

In principle, pABF can be combined with eABF, however, the mean force from eABF is already smoother than that from ABF,¹⁸ suggesting that the variance reduction benefit could be lower than with standard ABF.

5 Conclusion

I have presented a flexible and efficient multidimensional integration algorithm, which is integrated in the Colvars Module implementations of ABF and eABF-CZAR, so that they yield accurate free energy surfaces in dimensions 2 and 3. I have also implemented projected ABF, a reduced-variance version of ABF, together with new streamlined, efficient and flexible tools for visualization and analysis of ABF simulation results. Free energy gradients in dimension higher than 3, which are rarely used, can still be integrated using the existing integration tool in arbitrary dimension.⁹

I have also investigated pABF, which had previously been derived theoretically, and tested numerically on a toy model that, crucially, contained a single slow degree of freedom, without hidden barriers.¹¹ The tests show that pABF is a proven lower-variance method, with surprisingly deleterious consequences on exploration. These results are not sufficient to exclude a potential benefit of pABF in slower-relaxing systems, for which variance reduction and the associated slowdown in exploration could avoid long-lasting non-equilibrium biases in the free energy estimate.⁴⁰ However, it is possible that pABF reduces noise components with short autocorrelation times, not the long-lived biases due to slow orthogonal relaxation. Other modifications of ABF tend to reduce the diffusive timescales, such as stratification or combinations with metadynamics.³⁸ However, metadynamics and related methods, while producing faster exploration of collective variable space, do not address orthogonal barriers.

Stratification can be used to limit the diffusion times by reducing the space available to each trajectory. This strategy is compatible with pABF.

Comparing different methods on equal footing remains difficult, even when they are as close as ABF and pABF, because all tunable parameters should be optimized for each method independently. When some of the parameters play the role of adaptation rates, the presence of realistic relaxation, including a slow orthogonal space, is necessary to fully assess the real-world benefits of each method.

Acknowledgments

The author gratefully acknowledges support from the Laboratoire International Associé CNRS-UIUC, and from the French National Research Agency under grant LABEX DYNAMO (ANR-11-LABX-0011). The author thanks Dr. Tony Lelièvre for fruitful discussions, and Mr. Geoffrey Letessier for assistance with local computational resources. The author is grateful to the referees of this work for their constructive comments. Computational work was performed using HPC resources from LBT/HPC.

Supporting Information

The Supporting Information document contains: a description of the software architecture of the Colvars Module; a summary of the availability of the software, independently and within parent software packages; an example script using the provided Python module for analysis of ABF data.

This information is available free of charge via the Internet at <http://pubs.acs.org>

References

- (1) Huber, T.; Torda, A. E.; van Gunsteren, W. F. Local elevation: A method for improving the searching properties of molecular dynamics simulation. *J. Comput. Aided Mol. Des.* **1994**, *8*, 695–708.
- (2) Laio, A.; Parrinello, M. Escaping free energy minima. *Proc. Natl. Acad. Sci. U.S.A.* **2002**, *99*, 12562–12565.
- (3) Marsili, S.; Barducci, A.; Chelli, R.; Procacci, P.; Schettino, V. Self-healing umbrella sampling: a non-equilibrium approach for quantitative free energy calculations. *J. Phys. Chem. B* **2006**, *110*, 14011–14013.

- (4) Babin, V.; Roland, C.; Sagui, C. Adaptively biased molecular dynamics for free energy calculations. *The Journal of Chemical Physics* **2008**, *128*, 134101.
- (5) Dickson, B. M.; Legoll, F.; Lelièvre, T.; Stoltz, G.; Fleurat-Lessard, P. Free energy calculations: an efficient adaptive biasing potential method. *J. Phys. Chem. B* **2010**, *114*, 5823–5830.
- (6) Darve, E.; Pohorille, A. Calculating free energies using average force. *J. Chem. Phys.* **2001**, *115*, 9169–9183.
- (7) Kirkwood, J. G. Statistical mechanics of fluid mixtures. *J. Chem. Phys.* **1935**, *3*, 300–313.
- (8) Darve, E.; Rodríguez-Gómez, D.; Pohorille, A. Adaptive biasing force method for scalar and vector free energy calculations. *J. Chem. Phys.* **2008**, *128*, 144120.
- (9) Hénin, J.; Fiorin, G.; Chipot, C.; Klein, M. L. Exploring multidimensional free energy landscapes using time-dependent biases on collective variables. *J. Chem. Theory Comput.* **2010**, *6*, 35–47.
- (10) Lelièvre, T.; Rousset, M.; Stoltz, G. *Free Energy Computations: A Mathematical Perspective*; Imperial College Press, 2010.
- (11) Alrachid, H.; Lelièvre, T. Long-time convergence of an adaptive biasing force method: Variance reduction by Helmholtz projection. *SMAI-Journal of Computational Mathematics* **2015**, *1*, 55–82.
- (12) Fiorin, G.; Klein, M. L.; Hénin, J. Using collective variables to drive molecular dynamics simulations. *Mol. Phys.* **2013**, *111*, 3345–3362.
- (13) Hénin, J.; Chipot, C. Overcoming free energy barriers using unconstrained molecular dynamics simulations. *J. Chem. Phys.* **2004**, *121*, 2904–2914.

- (14) Ciccotti, G.; Kapral, R.; Vanden-Eijnden, E. Blue moon sampling, vectorial reaction coordinates, and unbiased constrained dynamics. *ChemPhysChem* **2005**, *6*, 1809–1814.
- (15) Cuendet, M. A.; Tuckerman, M. E. Free Energy Reconstruction from Metadynamics or Adiabatic Free Energy Dynamics Simulations. *Journal of Chemical Theory and Computation* **2014**, *10*, 2975–2986.
- (16) Mones, L.; Bernstein, N.; Csányi, G. Exploration, Sampling, And Reconstruction of Free Energy Surfaces with Gaussian Process Regression. *Journal of Chemical Theory and Computation* **2016**, *12*, 5100–5110.
- (17) Lelièvre, T.; Rousset, M.; Stoltz, G. Computation of free energy profiles with parallel adaptive dynamics. *J. Chem. Phys.* **2007**, *126*, 134111.
- (18) Lesage, A.; Lelièvre, T.; Stoltz, G.; Hénin, J. Smoothed Biasing Forces Yield Unbiased Free Energies with the Extended-System Adaptive Biasing Force Method. *J. Phys. Chem. B* **2017**, *121*, 3676–3685.
- (19) Kastner, J.; Thiel, W. Bridging the gap between thermodynamic integration and umbrella sampling provides a novel analysis method: “Umbrella integration”. *J. Chem. Phys.* **2005**, *123*, 144104.
- (20) Zheng, L.; Chen, M.; Yang, W. Simultaneous escaping of explicit and hidden free energy barriers: Application of the orthogonal space random walk strategy in generalized ensemble based conformational sampling. *J. Chem. Phys.* **2009**, *130*, 234105.
- (21) Fu, H.; Shao, X.; Chipot, C.; Cai, W. Extended Adaptive Biasing Force algorithm. An on-the-fly implementation for accurate free-energy calculations. *J. Chem. Theor. Comput.* **2016**, *12*, 3506–3513.
- (22) Kästner, J. Umbrella integration in two or more reaction coordinates. *The Journal of Chemical Physics* **2009**, *131*, 034109.

- (23) Abrams, C. F.; Vanden-Eijnden, E. On-the-fly free energy parameterization via temperature accelerated molecular dynamics. *Chemical Physics Letters* **2012**, *547*, 114–119.
- (24) Paz, S. A.; Maragliano, L.; Abrams, C. F. Effect of Intercalated Water on Potassium Ion Transport through Kv1.2 Channels Studied via On-the-Fly Free-Energy Parametrization. *Journal of Chemical Theory and Computation* **2018**, *14*, 2743–2750.
- (25) Maragliano, L.; Vanden-Eijnden, E. Single-sweep methods for free energy calculations. *J. Chem. Phys.* **2008**, *128*, 184110.
- (26) Stecher, T.; Bernstein, N.; Csányi, G. Free Energy Surface Reconstruction from Umbrella Samples Using Gaussian Process Regression. *Journal of Chemical Theory and Computation* **2014**, *10*, 4079–4097.
- (27) Shewchuk, J. R. *An Introduction to the Conjugate Gradient Method Without the Agonizing Pain*; School of Computer Science, Carnegie Mellon University, 1994.
- (28) Press, W.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P. *Numerical recipes in C++ : the art of scientific computing*; Cambridge University Press: Cambridge, UK New York, 2002; Section 2.7.
- (29) Lee, B.; Min, C. Optimal Preconditioners on Solving the Poisson Equation with Neumann Boundary Conditions. *433*, 110189.
- (30) Thomas, J. W. *Numerical partial differential equations: finite difference methods*; Springer Science & Business Media, 2013; Vol. 22.
- (31) Chen, L. Finite Difference Methods for Poisson Equation. Lecture Notes, 2019; <https://www.math.uci.edu/~chenlong/226/FDM.pdf>.
- (32) Phillips, J.; Hardy, D.; Maia, J.; Stone, J.; Ribeiro, J.; Bernardi, R.; Buch, R.; Fiorin, G.; Hénin, J.; Jiang, W.; McGreevy, R.; Melo, M. C. d. R.; Radak, B.; Skeel, R.; Singharoy, A.; Wang, Y.; Roux, B.; Aksimentiev, A.; Luthey-Schulten, Z.; Kale, L.;

- Schulten, K.; Chipot, C.; Tajkhorshid, E. Scalable molecular dynamics on CPU and GPU architectures with NAMD. *The Journal of Chemical Physics* **2020**, *153*, 044130.
- (33) Barducci, A.; Bussi, G.; Parrinello, M. Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Phys. Rev. Lett.* **2008**, *100*, 020603.
- (34) Altis, A.; Nguyen, P. H.; Hegger, R.; Stock, G. Dihedral angle principal component analysis of molecular dynamics simulations. *J. Chem. Phys.* **2007**, *126*, 244111.
- (35) Glykos, N. M. Carma: a molecular dynamics analysis program. *J. Comput. Chem.* **2006**, *27*, 1765–1768.
- (36) Humphrey, W.; Dalke, A.; Schulten, K. VMD: visual molecular dynamics. *J. Mol. Graph.* **1996**, *14*, 33–8, 27–8.
- (37) Paz, S.; Abrams, C. Testing Convergence of Different Free-Energy Methods in a Simple Analytical System with Hidden Barriers. *Computation* **2018**, *6*, 27.
- (38) Fu, H.; Shao, X.; Cai, W.; Chipot, C. Taming Rugged Free Energy Landscapes Using an Average Force. *Accounts of Chemical Research* **2019**, *52*, 3254–3264.
- (39) Zheng, L.; Chen, M.; Yang, W. Random walk in orthogonal space to achieve efficient free-energy simulation of complex systems. *Proc. Natl. Acad. Sci. U.S.A.* **2008**, *105*, 20227–20232.
- (40) Miao, M.; Fu, H.; Zhang, H.; Shao, X.; Chipot, C.; Cai, W. Avoiding non-equilibrium effects in adaptive biasing force calculations. *Molecular Simulation* **2020**, 1–5.

For Table of Contents Only

