

Exploration, sampling and reconstruction of free energy surfaces with Gaussian process regression

Letif Mones,^{*,†} Noam Bernstein,[‡] and Gábor Csányi[†]

[†]*Engineering Laboratory, University of Cambridge, Cambridge, CB2 1PZ, United Kingdom*

[‡]*Naval Research Laboratory, Center for Materials Physics and Technology, Washington, DC 20375, United States of America*

E-mail: lam81@cam.ac.uk

Abstract

Practical free energy reconstruction algorithms involve three separate tasks: biasing, measuring some observable, and finally reconstructing the free energy surface from those measurements. In more than one dimension, adaptive schemes make it possible to explore only relatively low lying regions of the landscape by progressively building up the bias towards the negative of the free energy surface so that free energy barriers are eliminated. Most schemes use the final bias as their best estimate of the free energy surface. We show that large gains in computational efficiency, as measured by the reduction of time to solution, can be obtained by separating the bias used for dynamics from the final free energy reconstruction itself. We find that biasing with metadynamics, measuring a free energy gradient estimator, and reconstructing using Gaussian process regression can give an order of magnitude reduction in computational cost.

1 Introduction

Computing free energy differences and entire free energy surfaces of complex systems that have a large number of intrinsic degrees of freedom plays a central role in computational chemistry and biochemistry. Reducing the high dimensional potential energy surface to a much lower dimensional free energy landscape helps to identify and focus on the most relevant configurations of the system and the transitions between them, and can be used both for understanding chemical processes and for validating computational models.

Although there are methods to obtain absolute free energies by computing the partition function directly, such as Wang-Landau sampling¹ and Nested Sampling,² these are rather expensive, and are only just beginning to be used for chemical systems and even then are competitive only in cases where many thermodynamic functions need to be computed simultaneously in a wide temperature range. More typically, free energy differences between specific states at a given temperature can be obtained by simulation techniques such as free energy perturbation³ and thermodynamic integration.^{4,5} These methods are widely used to calculate solvation, hydration, and binding free energies.⁶

However, when a more detailed "mechanistic" view is required, it is customary to introduce a number of *collective variables* to represent relevant degrees of freedom, which are often the dynamically slower ones, and calculate free energy surfaces in this collective variable space. The free energy A at a particular collective variable value ξ^* is defined as the logarithm of the marginal probability density,

$$A(\xi^*) = -\beta^{-1} \ln P(\xi^*) = -\beta^{-1} \ln \int \delta(\xi(\mathbf{r}) - \xi^*) e^{-\beta V(\mathbf{r})} d\mathbf{r} \quad (1)$$

where $\xi(\mathbf{r})$ is the set of collective variables represented as a vector function of the atomic coordinates \mathbf{r} , $P(\xi^*)$ is the probability density at ξ^* , β is the inverse temperature, and we suppress the momentum dependence which leads only to an additive constant. Such a free energy function is often called the potential of mean force (PMF), because it can also be calculated via its gradient,

which has a particularly simple estimator if the collective variables are explicit variables of the Hamiltonian.⁴ For instance, for $\xi(\mathbf{r}) \equiv \mathbf{r}_s$, where the components of \mathbf{r}_s are an arbitrary subset of the components of \mathbf{r} :

$$\nabla_{\mathbf{r}_s} A(\mathbf{r}_s^*) = \langle \nabla_{\mathbf{r}_s} V(\mathbf{r}) \rangle_{\mathbf{r}_s = \mathbf{r}_s^*} \quad (2)$$

where the term inside the ensemble average is the negative of the force components in the directions corresponding to the collective variable.

Another reason to introduce collective variables is to increase computational efficiency in the presence of free energy barriers. For atomistic models of complex systems, the practically feasible simulation time using molecular dynamics (MD) ranges from nano- to milliseconds depending on computational capacity and the size of the system. This time scale is often many orders of magnitude less than the intrinsic timescale of chemical processes. The common solution is to bias the dynamics along the intrinsically slow degrees of freedom. This can be done by selecting these slow degrees of freedom as collective variables, and using a bias that is expressed as an additional energy or force dependent on the collective variable values. Note that collective variables selected for the purpose of speeding up the exploration may be different from collective variables of interest for their mechanistic interpretation. The essential role of the bias is to help to sample configurations that otherwise would not be visited under unbiased dynamics within reasonable time. For simple chemical reactions there are good general guidelines for selecting reaction coordinates,⁷⁻⁹ but for more complex cases it can often be tricky to guess what the good collective variables are, and there has been significant effort to try and deduce them directly from simulations.^{10,11} A simpler alternative is to increase the number of collective variables in the hope that they span the subspace containing the slow degrees of freedom, but increasing the dimensionality of the collective variable space also significantly increases the cost of exploring it and the complexity of strategies for efficient exploration.

A key aspect of the free energy methods we discuss below is whether they use eq 1 or eq 2 to estimate the free energy surface. In the former case, the estimate depends on the relative frequency of visits to different parts of collective variable space, and therefore many revisits are required to

achieve convergence within sensible tolerances. In contrast, when using eq 2, the measurement of the free energy gradient is independent at each point in collective variable space, and its convergence is only dependent on the variance of the estimator and the number of visits at a given collective variable location, and not the relative frequency of visits at different locations. Using eq 1 is like solving a probability density estimation problem whereas using eq 2 leads to a regression problem. The question of which of these inference tasks is harder depends on the decorrelation time scales of the sampling process and variance of the gradient estimator, which we discuss below.

In general one is only interested in the relatively low free energy regions of the full collective variable space. A priori it is not known where these regions are, so an adaptive bias that preferentially enhances sampling in such regions is required. A common way to achieve this is to progressively update the bias during the simulation to the negative of the estimate of the free energy surface using the data measured thus far, using eq 1 or eq 2. It is natural to use the negative of the final bias as the best estimate of the free energy surface. However, setting the bias and reporting the final free energy surface are not intrinsically linked. For example, Chipot and Lelièvre introduced an efficient adaptive biasing force (ABF) method variant for weakly coupled collective variables,⁹ where the biasing force is composed of one-dimensional functions that enhance the sampling in the multidimensional collective variable space, and proposed combining it with a standard computation of the multidimensional free energy gradients for reconstruction of the free energy surface. A similar division of tasks was recently proposed by Fiorin et al.,¹² who suggested using metadynamics as the bias while computing the free energy profile from thermodynamic integration of the gradient estimator provided by the ABF method. They applied this combination of approaches to a simple one dimensional system, but did not report its computational cost. Another approach was introduced by Cuendet and Tuckerman,¹³ who used an extended Lagrangian formalism and a gradient estimator based on the derivative conjugate to the extended degree of freedom. However, this combination did not result in a significant gain in efficiency compared to conventional adaptive methods.

In this work we show that there is a substantial advantage to using different estimators for the

biasing and reporting tasks, rather than using the same estimate for both. We begin with a toy model of a two dimensional Gaussian potential energy surface. We show how the performances of the density and gradient estimators vary as a function of the curvatures of the potential energy surface along its slow and fast directions and the orientation of the slow direction with respect to the collective coordinate. We then test the performance of the metadynamics^{14,15} (MTD) kernel density estimator and the adaptive biasing force (ABF) method¹⁶ gradient estimator for exploration and reconstruction on the alanine dipeptide and tripeptide in the gas phase and in solution. We find the best performance, i.e. lowest error in the reconstructed free energy surface for fixed amount of work, by using metadynamics to bias the exploration while sampling the gradient like ABF and reconstructing the final free energy surface from eq 2 using Gaussian process regression (GPR).¹⁷

2 Background

In biased dynamics the biasing force \mathbf{F}_b is used to modify the trajectory via an additive term in the definition of the effective force that is used in the equations of motion

$$\mathbf{F}_{\text{eff}} = -\nabla_{\mathbf{r}}V(\mathbf{r}) + \nabla_{\mathbf{r}}\xi(\mathbf{r}) \cdot \mathbf{F}_b(\xi, t). \quad (3)$$

There are two different approaches for what the bias should accomplish. In the first, the bias is used to tightly restrain the region of collective variable space that the system visits, and this region is explicitly scanned over all the interesting parts of the space, e.g. near relevant minima and barriers that separate them. Because we are interested in the situation where these regions are not known *a priori*, we focus on the alternative approach, in which the bias is used to artificially speed up exploration. Without the bias, the probability of finding the system at a particular collective variable position scales inversely exponentially with the corresponding free energy. If the PMF was known, its negative could be added as a biasing potential $V_b(\xi, t)$ (with the biasing force defined as the negative of its gradient: $\mathbf{F}_b(\xi, t) = -\nabla_{\xi}V_b(\xi, t)$) and the system would see a “flat” free energy landscape and move purely diffusively. This is the idea behind most adaptive biasing

schemes: to build up an ever-more-accurate picture of the free energy surface during a simulation, and use its negative as an ever-changing bias of the same simulation. Careful attention has to be paid to the consistency of such schemes and some equilibrium must be reached in the limit $t \rightarrow \infty$ (i.e. $V_b(\boldsymbol{\xi}, t) \rightarrow V_b(\boldsymbol{\xi}, \infty)$ and $\mathbf{F}_b(\boldsymbol{\xi}, t) \rightarrow \mathbf{F}_b(\boldsymbol{\xi}, \infty)$ for adaptive biasing potential and adaptive biasing force methods, respectively).^{18,19}

One approach is to estimate the free energy surface at each time from the trajectory up to that time point by estimating the integral in eq 1 from the histogram of occupancies of bins in collective variable space. This is what is done in the self-healing umbrella sampling,²⁰ adaptive biasing potential,²¹ and local elevation umbrella sampling²² methods. As mentioned before, the convergence of the histogram requires that the system transitions between bins many times to correctly estimate their relative probabilities. Histogram-based estimates of probability densities become very noisy in higher dimensions due to the exponential growth of the number of bins and the consequent low bin counts. Methods like local elevation,²³ conformational flooding,²⁴ and metadynamics^{14,15} alleviate this latter problem to some extent by using kernel density estimators. For example, in the case of MTD the probability density is approximated as the product of log-normal distributions. This is equivalent to the conventional view that metadynamics approximates the free energy surface as a sum of Gaussians.

A method that is different from those above in many respects is the adaptive biasing force technique.¹⁶ Instead of constructing a global estimate of the free energy surface to be used as a biasing potential, ABF uses the trajectory data to estimate the *gradient* of it directly and pointwise (i.e. as a conditional average at each collective variable location), and uses this as a biasing force \mathbf{F}_b for constructing the effective force in eq. 3. Traditionally, the conditional average is computed in a discretised way using bins with fixed size¹⁶ but a continuous variant is also possible using a kernel estimator.¹⁸ The statistical estimator for \mathbf{F}_b is not necessarily simple or even unique (see below); the only requirement is that its average should approach $\nabla_{\boldsymbol{\xi}} A(\boldsymbol{\xi}^*)$ as the trajectory evolves. In contrast to the previously mentioned adaptive methods, ABF does not necessarily require many revisits in collective variable space, because the convergence of the gradient estimator is independent

at each location, and is assured without the need for any further empirical tuning.^{18,25} However, using ABF to reconstruct the free energy surface for more than one collective variable requires a non-trivial multidimensional integration of the sampled gradient.²⁶ Here we accomplish this using Gaussian process regression, which in our previous work¹⁷ was shown to be an effective and robust technique for reconstructing the complete free energy surface from a set of fixed bias simulations, e.g. umbrella sampling of small peptides. In that case the advantage over alternative methods tested was marginal in one dimension, significant in two dimensions, and overwhelming in four dimensions.

3 Methods

In this section we briefly review the formalism of free energy gradient estimators and show how they can be used with Gaussian process regression. We then describe the other free energy surface methods we compare our results with, our test systems, and the computational details.

3.1 Instantaneous collective forces from unbiased and biased simulations

For arbitrary collective variables eq 2 can be generalized by taking the gradient of eq 1 and carrying out some additional manipulations^{18,27} to express the free energy gradient as an ensemble average of a force-like quantity \mathbf{f}

$$\nabla_{\xi} A_u(\xi^*) = -\langle \mathbf{f} \rangle_{\xi=\xi^*, u} \quad (4)$$

where the dimension of \mathbf{f} is D , the number of collective variables, the subscript “u” indicates that the quantity is calculated from an unbiased simulation, and $\langle \dots \rangle_{\xi^*, u}$ denotes an unbiased conditional average at a collective variable value of ξ^* . We refer to \mathbf{f} as the *instantaneous collective force* (ICF), because it is a sample at a single point in time of the force conjugate to a collective

variable. It is sometimes also called the local mean force.¹⁸ The general form of ICF is¹⁸

$$\mathbf{f} = - (\mathbf{G}_W^{-1} \mathbf{W}) \cdot \nabla_{\mathbf{r}} V + \beta^{-1} \nabla_{\mathbf{r}} \cdot (\mathbf{G}_W^{-1} \mathbf{W}), \quad (5)$$

where $\mathbf{G}_W^{-1} \in \mathbb{R}^{D \times D}$ is the inverse of the generalised Gram matrix $\mathbf{G}_W = \mathbf{W} \nabla_{\mathbf{r}} \boldsymbol{\xi}$, the term $\nabla_{\mathbf{r}} \boldsymbol{\xi} \in \mathbb{R}^{3N \times D}$ is the transpose of the Jacobian matrix of the collective variables, and $\mathbf{W} \in \mathbb{R}^{D \times 3N}$ is an arbitrary vector field with the restriction that \mathbf{G}_W must be invertible. For a specific problem an appropriate choice of \mathbf{W} can significantly reduce the variance of \mathbf{f} resulting in faster convergence. The choice $\mathbf{W} = \nabla_{\mathbf{r}} \boldsymbol{\xi}^T \boldsymbol{\mu}^{-1}$, where $\boldsymbol{\mu} \in \mathbb{R}^{3N \times 3N}$ is the mass matrix with $\mu_{ij} = \delta_{ij} m_i$, leads to $\mathbf{G}_W = \nabla_{\mathbf{r}} \boldsymbol{\xi}^T \boldsymbol{\mu}^{-1} \nabla_{\mathbf{r}} \boldsymbol{\xi}$, which is usually denoted by \mathbf{Z} and is called the metric tensor.

With the special choice of $\mathbf{W} = \nabla_{\boldsymbol{\xi}} \mathbf{r}$ the generalised Gram matrix becomes the identity matrix and equation 5 can be written as:^{27,28}

$$\mathbf{f} = -\nabla_{\boldsymbol{\xi}} V + \beta^{-1} \nabla_{\boldsymbol{\xi}} \ln |\mathbf{J}|, \quad (6)$$

where $\mathbf{J} = \nabla_{\mathbf{q}} \mathbf{r}$ is the full Jacobian matrix (where \mathbf{q} is the set of generalised coordinates that includes both $\boldsymbol{\xi}$ and the coordinates in the perpendicular space), i.e. the gradient of the inverse transformation from the generalised coordinates to the Cartesian ones. Whatever vector field \mathbf{W} is chosen, the second part of eq 5 requires the calculation of the second derivatives of the collective variables, which can be cumbersome for some collective variables, e.g. dihedral angles. To overcome this difficulty, an elegant formula was derived by Darve et al.:^{16,29,30}

$$\mathbf{f} = \frac{d}{dt} \left(\mathbf{Z}^{-1} \frac{d\boldsymbol{\xi}}{dt} \right) - \mathbf{F}_b(\boldsymbol{\xi}) \quad (7)$$

where \mathbf{Z}^{-1} is the inverse of the metric tensor introduced before. The time derivatives in the first term in the right hand side can be approximated numerically by a finite difference, which eliminates the need for calculating the second derivatives and is the basis of the ABF method.¹⁶ This approximation is used in this work through its implementation in PMFlib.³¹ An alternative approach,

applying an appropriate transform of the Cartesian coordinates to generalised coordinates and using inverse gradients to avoid the explicit computation of the second derivatives in eq 6,^{28,32,33} is used by the NAMD package.³⁴

One of the most important advantages of using ICFs for free energy reconstruction is that if the applied bias depends solely on the collective variables (as is almost always the case) then the conditional average of ICFs is invariant to this bias:

$$\langle \mathbf{f} \rangle_{\xi=\xi^*,b} = \langle \mathbf{f} \rangle_{\xi=\xi^*,u} = -\nabla_{\xi} A_u(\xi^*) \quad (8)$$

where the subscripts “b” and “u” denote that the quantity is calculated from biased and unbiased simulations, respectively. For methods that do not define any biasing potential but directly apply a biasing force (like ABF), one could in principle introduce a constraint that the biasing vector field is a gradient of some potential and in this case the equation formally holds for these methods as well.

Equation 8 only states that if the dynamics is at *equilibrium* then the average of the observed ICFs from a biased simulation can be used as the estimate of the free energy gradient. For methods that apply a constant bias, equilibrium is eventually reached. However, for adaptive methods the bias is time dependent, which in general prevents the existence of true equilibrium. In practice and with some care one can maintain quasi-equilibrium even for adaptive methods (i.e. the change in the bias is slower than the dynamics on the underlying potential, which is automatically satisfied once the bias is converged).

3.2 Gaussian process regression

Gaussian process regression (GPR) is a non-parametric machine learning technique that can be used as a tool to reconstruct functions in multiple dimensions from noisy observations of their values and/or derivatives in a way that is physically motivated, robust, and free of overfitting artefacts. The theory can be found in detail in Refs. 35 and 36, and we have shown previously how to apply

it to free energy reconstruction.¹⁷ Here we just review the formulas that are used for this work. A covariance of function values is defined in the input space (in our case the space of collective variables), which sets the form of the basis functions, and can be used to impart smoothness to the reconstructed free energy via a length scale parameter. For example one can use a squared exponential covariance kernel

$$K_{\text{SE}}(A(\boldsymbol{\xi}^{(i)}), A(\boldsymbol{\xi}^{(j)})) = \chi^2 \exp \left(-\frac{1}{2} \sum_{k=1}^D \frac{(\xi_k^{(i)} - \xi_k^{(j)})^2}{\theta_k^2} \right) \quad (9)$$

where the super- and subscripts on ξ refer to different locations and components of the collective variables. In the present work, however, the free energy is periodic in its collective variables, in which case it is worth building this into the covariance function by using a periodic squared exponential,

$$K_{\text{PSE}}(A(\boldsymbol{\xi}^{(i)}), A(\boldsymbol{\xi}^{(j)})) = \chi^2 \exp \left(-2 \sum_{k=1}^D \frac{\sin^2 \left(\frac{\xi_k^{(i)} - \xi_k^{(j)}}{\theta_k} \right)}{\theta_k^2} \right) \quad (10)$$

The covariance function has a few free parameters (called hyperparameters in the machine learning literature), which are fixed using chemical understanding: $\boldsymbol{\theta}$ is a vector of typical length scales of features of the free energy landscape, and χ^2 is the expected variance of the overall surface. The reconstruction is fairly robust to the variation of these parameters within a factor of two or so.¹⁷

In our case we use only derivative observations, and the reconstruction takes the form

$$A(\boldsymbol{\xi}^*) = -\mathbf{k}(\boldsymbol{\xi}^*)^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{f} \quad (11)$$

where \mathbf{f} is a vector of derivative observations and σ represent their associated error (and is another hyperparameter), and

$$[\mathbf{k}(\boldsymbol{\xi}^*)]_i = \nabla_{\boldsymbol{\xi}} K(A(\boldsymbol{\xi}^{(i)}), A(\boldsymbol{\xi}^*)) \quad (12)$$

$$[\mathbf{K}]_{ij} = \nabla_{\boldsymbol{\xi}} \nabla_{\boldsymbol{\xi}'} K(A(\boldsymbol{\xi}^{(i)}), A(\boldsymbol{\xi}'^{(j)})) \quad (13)$$

where ξ^* is the collective variable location where we wish to predict the free energy. In the one dimensional case i and j are indices of previously observed free energy gradients, and in higher dimensions they refer to blocks.¹⁷ The reconstruction formula, eq 11, can be understood as a linear superposition of basis functions given by the covariance kernel derivatives, with coefficients determined by the observations of gradients and their mutual covariance.

3.3 Free energy surface reconstruction from instantaneous collective forces

Let us suppose that we can afford a fixed number of MD steps, and we want to calculate the free energy surface from derivative observations using GPR. The following question arises: how should the observations be distributed in collective variable space? If a grid is used, there is a trade-off between the number of grid points and the number of MD steps per grid point: using fewer grid points allows a larger number of samples per grid point, making the derivative estimate more precise (because the error on the mean scales with the inverse of the square root of the number of samples), but at the same time might result in larger error in the integrated surface due to the larger grid spacing. Our earlier investigation of Gaussian process regression free energy surface reconstruction using umbrella sampling¹⁷ showed that at fixed total sampling time, a larger number of umbrella centers with less sampling at each is better than a smaller number of umbrella centers with more sampling at each, despite the larger noise in the sampled gradients in the former. By fitting a scalar free energy function the regression process automatically removes components of the noise with non-zero curl, which would corrupt direct integration along specific paths (e.g. thermodynamic integration) in a multidimensional noisy sampled gradient field.

In this work we go to the extreme limit of this trade-off: without doing any local averaging, we use the collected ICFs directly according to eq 4 as noisy estimates of the free energy derivatives. In some sense this corresponds to one sample per “grid point” although we do not construct a grid explicitly, just note the collective variable location of each derivative observation. So rather than first estimating the gradient as an average of the ICFs and then carrying out the integration, we combine the two steps.

The most computationally expensive part of GPR is the matrix inversion in eq 11 and therefore the method scales as $O(N^3D^3)$. For a large data set, e.g. the millions of data points that would arise from sampling gradients from a long MD run, this is not feasible. The computational cost can be reduced by using sparsification techniques³⁷ that project the N observations to a smaller number of “sparse points” $M \ll N$ that represent the entire data set. The sparse points can (but do not have to) be a subset of the original data set. The computation then involves smaller covariance matrices between the set of all observed data and the set of sparse points (size $ND \times MD$) and among the sparse points themselves (size $MD \times MD$). The sparse method scales as $O(M^2ND^3)$. The sparse points can be chosen on a predefined grid, but for higher dimensional problems, where the exploration of the entire landscape is not needed, clustering methods such as k-means³⁵ or CUR³⁸ can be used.

3.4 Summary of free energy surface methods used

The method we introduce in this paper uses metadynamics (MTD) to bias the motion of the system while sampling instantaneous collective forces (ICF) that are used to reconstruct the final free energy surface with Gaussian process regression (GPR), a combination we refer to as MTD/ICF/GPR. In addition to using previously published PMF methods (or minor variations thereof) as components in our new scheme, we also present results from these and other methods for comparison purposes and to generate reference data. A detailed description of the actual parameters used for the different methods can be found in the supporting information (SI). One method is metadynamics,¹⁵ which uses a biasing potential that consists of a sum of repulsive hills periodically deposited at previous locations along the trajectory in collective variable space. The hills are multidimensional Gaussians in the Euclidian distance in collective variable space with fixed widths, and heights that decrease during the simulation according to the well-tempered MTD scheme.¹⁹ Final reconstruction of the free energy surface is obtained by averaging the biasing potential over a range of simulation time.⁷ Another is the adaptive biasing force method,¹⁶ which uses a biasing force from sampled ICFs that are averaged in bins. For the multidimensional inte-

gration necessary for the final free energy surface reconstruction we use GPR on the final averaged biasing forces.

We also test ABF(GPR), a binless variant of ABF where we use the sampled ICFs directly in the GPR to generate the biasing force. This makes ABF(GPR) an example of ABF with a kernel-based bias force reconstruction,¹⁸ albeit with an unusual choice of kernels for the bias force, namely the gradients of Gaussian functions. Since the biasing forces need to be updated during the simulation to include new samples, here sparsification is essential to reduce the computational cost dominated by covariance matrix inversions, and the updating of the bias potential is only done every one thousand MD time steps.

To generate reference free energy surfaces we use blue moon (BM) sampling.³⁹ In this non-adaptive method the system is kept at a specific collective variable value using constrained equations of motion, and the free energy gradient is estimated from the sampled Lagrange multipliers corresponding to the constraints. To generate the entire surface, gradients need to be sampled at a grid of collective variable values. Note that generating the initial configuration for the simulation at each grid point is not trivial: the system must first be equilibrated, and this can be a significant computational cost. We start from an equilibrated configuration from a nearby collective variable value, making the sampling process sequential over grid points.

In the “flat potential” method we use the reference PMF to perfectly bias the dynamics, so that there are no free energy barriers and exploration is purely diffusive. This gives a limiting case for the exploration rate which skips the task of building up a good bias, and helps us characterise the rate of diffusion in collective variable space.

3.5 Description of the model systems

3.5.1 Harmonic oscillator systems

We start our investigation with a comparison between the density and gradient-based free energy surface reconstructions in a simple toy model. The two reconstruction strategies were used for a 2-dimensional harmonic oscillator model system, where we have chosen one of the two co-

ordinates as the collective variable.

The potential energy function of the system is defined as the quadratic form

$$V_{\sigma^2, \phi}(x, y) = \frac{1}{2} \begin{pmatrix} x & y \end{pmatrix} \mathbf{C}_{\sigma^2, \phi}^{-1} \begin{pmatrix} x \\ y \end{pmatrix} \quad (14)$$

where

$$\mathbf{C}_{\sigma^2, \phi}^{-1} = \mathbf{R}_\phi \mathbf{C}_{\sigma^2}^{-1} \mathbf{R}_\phi^{-1} \quad (15)$$

is the matrix of inverse principal curvatures of the potential

$$\mathbf{C}_{\sigma^2}^{-1} = \begin{pmatrix} 1/\sigma_1^2 & 0 \\ 0 & 1/\sigma_2^2 \end{pmatrix} \quad (16)$$

rotated by an angle ϕ via the rotation matrix \mathbf{R}_ϕ . Defining $c = \cos \phi$ and $s = \sin \phi$, the explicit form of eq 14 becomes

$$V_{\sigma^2, \phi}(x, y) = \frac{1}{2} \left(\left[\frac{1}{\sigma_1^2} c^2 + \frac{1}{\sigma_2^2} s^2 \right] x^2 + 2 \left[\frac{1}{\sigma_1^2} cs - \frac{1}{\sigma_2^2} cs \right] xy + \left[\frac{1}{\sigma_1^2} s^2 + \frac{1}{\sigma_2^2} c^2 \right] y^2 \right) \quad (17)$$

The canonical probability density function of the system is therefore a normal distribution:

$$P_{\sigma^2, \phi, \beta}(x, y) = \mathcal{N} \left(0, \beta^{-1} \mathbf{C}_{\sigma^2, \phi} \right) = \frac{1}{\sqrt{(2\pi)^2 \beta^{-2} \sigma_1^2 \sigma_2^2}} e^{-\beta V_{\sigma^2, \phi}(x, y)} \quad (18)$$

Now we select the coordinate x as the collective variable along which we want to compute the free energy. The exact free energy profile can be easily obtained by marginalising the joint probability distribution above:

$$\begin{aligned} P_{\sigma^2, \phi, \beta}(x) &= \mathcal{N} \left(0, \beta^{-1} [\mathbf{C}_{\sigma^2, \phi}]_{11} \right) = \mathcal{N} \left(0, \beta^{-1} (\sigma_1^2 c^2 + \sigma_2^2 s^2) \right) \\ &= \frac{1}{\sqrt{2\pi \beta^{-1} (\sigma_1^2 c^2 + \sigma_2^2 s^2)}} e^{-\frac{1}{2} \beta (\sigma_1^2 c^2 + \sigma_2^2 s^2)^{-1} x^2} \end{aligned} \quad (19)$$

The free energy profile along x is by definition the negative logarithm of the marginal distribution:

$$A_{\sigma^2, \phi, \beta}(x) = -\beta^{-1} \ln P_{\sigma^2, \phi, \beta}(x) = -\beta^{-1} \ln \frac{1}{\sqrt{2\pi\beta^{-1}(\sigma_1^2 c^2 + \sigma_2^2 s^2)}} + \frac{1}{2} (\sigma_1^2 c^2 + \sigma_2^2 s^2)^{-1} x^2 \quad (20)$$

Note that this general equation includes only 3 free parameters (σ_1^2 , σ_2^2 and ϕ) and apart from a constant the free energy profile does not depend on temperature. Finally, the ICF corresponding to the collective variable x is simply equal to the derivative of the potential energy function:

$$f_x = -\frac{\partial V_{\sigma^2, \phi}}{\partial x} = -\left[\frac{1}{\sigma_1^2} c^2 + \frac{1}{\sigma_2^2} s^2\right] x - \left[\frac{1}{\sigma_1^2} cs - \frac{1}{\sigma_2^2} cs\right] y \quad (21)$$

3.5.2 Molecular systems

Four molecular model systems were investigated: alanine dipeptide (*N*-acetyl–alanine–*N*'-methylamide, ACE–ALA–NME) and alanine tripeptide (*N*-acetyl–alanine–alanine–*N*'-methylamide, ACE–ALA–ALA–NME) in gas phase and water solution. The peptides were described by the Amber ff99SB force field.⁴⁰ For the simulations in water solution we used the flexible TIP3P⁴¹ potential and periodic boundary conditions. Long range non–bonded interactions were computed by the particle mesh Ewald method⁴² with a direct interaction cutoff of 9 Å. The collective variables were the dihedral angles: $\Phi(\text{C} - \text{N} - \text{C}_\alpha - \text{C})$ and $\Psi(\text{N} - \text{C}_\alpha - \text{C} - \text{N})$, leading to 2 dimensional free energy surface for alanine dipeptide and 4 dimensional ones for alanine tripeptide (for details see Figure 1). Simulations were performed using the Amber14 package⁴³ linked to PMFlib,³¹ which was used for the calculation of the biases and free energy surface estimates.

We investigated two properties as functions of the total simulation time: the rate of the exploration of the entire free energy landscape and the root mean squared (RMS) error of the multidimensional PMF estimate. To measure the exploration rate we divided the entire collective variable space into bins (48^2 for 2-D and 12^4 for 4-D) and a given bin was considered visited if the system entered the bin at least once during the simulation. For additional computational details see the SI.

4 Results and Discussion

4.1 Harmonic oscillator systems

Figure 2 illustrates the probability distribution of $x - y$ (first row) and the marginal distribution along x (second row) at fixed variances ($\sigma_1^2 = 1.0$ and $\sigma_2^2 = 0.04$), $\beta = 1$ and three different ϕ values using 1 million i.i.d. samples. We also show the distribution of the ICF (f_x) for 3 values of x (third row). It is clear from the figure that increasing the coupling between the variables (i.e. rotating the distribution anti-clockwise) makes the ICF distribution wider and so the estimated gradients will have an increased error. This observation is in accord with the analytical variance of the ICFs

$$\text{var}(f_{\sigma^2, \phi} | x = x^*) = \beta^{-1} \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_2^2} \right)^2 \frac{c^2 s^2 \sigma_1^2 \sigma_2^2}{\sigma_1^2 c^2 + \sigma_2^2 s^2} \quad (22)$$

Note that this variance does not depend on x and is exactly zero at $\phi = 0$ and $\pi/2$.

To evaluate the error of the free energy surface as a function of dimensionless potential energy parameters σ_1/σ_2 and ϕ , we numerically sampled gradients and histograms (additional details in the SI). We measure the RMS of the free energy surface error over a range of collective variable values proportional to the width of the free energy distribution. In Figure 3 we compare the results for the two methods. While the error of the histogram-based method is constant, the error of the ICF-based reconstruction is sensitive to the model parameters. The latter has better overall performance when the curvatures of the two variables do not differ by more than 2 orders of magnitude or ϕ is close to 0 or $\pi/2$.

We note that for a realistic system, the error in an ICF-based reconstruction will not be controlled by a single curvature. Nevertheless it will depend on the variance of the ICF over the space perpendicular to the collective variables, while for histogram-based reconstructions, the error is controlled by diffusion rates in collective variable space.

4.2 Alanine peptide exploration rate

The explored fractions of the collective variable volume (fraction of bins sampled at least once) for the alanine dipeptide and tripeptide in gas phase and in solution as a function of simulation time are plotted in Figure 4 and Figure 5. The unbiased simulations quickly explore a small fraction of the space but then stagnate, unable to cross the free energy barriers. The flat potential method does much better, exploring the entire space within our simulation times. While it might be naively expected that a flat potential would be optimal for exploration, MTD actually explores collective variable space more quickly for both peptides in solution, and for the dipeptide in the gas phase. Note that choosing the parameters of the MTD repulsive hills and well-tempered algorithm requires a tradeoff between fast exploration (large frequent hills) and minimal deviation from equilibrium (small infrequent hills). Similarly, in the case of ABF methods the exploration rate is a function of the number samples collected before applying the ABF force. We chose the parameters to reduce the overall reconstruction error as quickly as possible. Accordingly, for MTD the initial heights of the hills were 1.0 and 10.0 kcal mol⁻¹ for the alanine dipeptide and tripeptide, respectively. The widths of the hills were $\pi/10$ rad in all directions and deposition intervals of 50 and 500 fs were applied for the dipeptide and tripeptide, respectively. For conventional ABF and ABF(GPR) methods we updated the bias at every 1 and 1000 steps, respectively. Additional details can be found in SI.

The ABF method explores the space much more slowly, although it does not become trapped like the unbiased simulation. Replacing the ABF biasing force with a GPR reconstructed one based on all gradient samples, i.e. using ABF(GPR), greatly improves on conventional ABF in 2-D, matching the flat potential rate of exploration. It therefore appears that binning the sampled gradients into a piecewise-constant biasing force is a significant contribution to the relatively slow exploration rate of ABF and that using a kernel-based reconstruction of the bias force, as first proposed in Ref.,¹⁸ can lead to an improvement.

The fact that MTD leads to faster exploration compared to any of the ABF variants tested may be a consequence of the qualitatively different ways these methods modify the biasing forces dur-

ing the simulation. MTD always places repulsive hills in previously visited locations, leading to self avoiding behavior at least on short timescales. ABF, whether histogram or kernel based, compensates for local gradients, and therefore behaves differently depending on the local curvature. Near minima it creates bias forces away from the minimum, i.e. self-avoiding, near maxima it creates bias forces toward the maximum, i.e. reducing the tendency of the system to move away, and in between it produces bias forces that counter the local gradient.

Since MTD still outperforms ABF(GPR) for exploration, in the tripeptide system we did not test the ABF(GPR) method because the simulations are longer and the gradients have higher dimension, making the cost of repeatedly updating the GPR bias during the simulation quite high.

4.3 Alanine peptide free energy surface error

The RMS errors of the final reconstructed free energy surfaces for the alanine dipeptide and tripeptide in gas phase and in solution as a function of simulation time are plotted in Figure 6 and Figure 7. In addition to MTD, ABF, and MTD/ICF/GPR, we also plot two curves based on blue moon sampling with the same size grids (48^2 and 12^4) used to bin the ABF biasing forces. The one labelled BM(only-sampling) only takes into account the time spent sampling the Lagrange multipliers at each grid point. It ignores the simulation time needed to create an equilibrated initial configuration at each grid point, usually done by slowly moving the constraint from the previous grid point and then equilibrating at the new position. It is not a complete PMF method, but rather serves as a lower bound for the sampling time of gradient based reconstruction, only measuring the time for convergence of the gradient estimator. The one simply labelled BM includes the simulation time to move and equilibrate the initial configuration at each grid point, and is a complete free energy surface reconstruction method, although since it is not adaptive it is restricted to reconstructing the free energy surface over the entire collective variable domain. Alternatively, BM could be considered a reference method that artificially removes the task of determining *where* in collective variable space needs to be explored, but still includes the time to move to and equilibrate at grid points that span the entire space.

In general the MTD and ABF methods are able to explore and reconstruct the free energy surface, although their relative performance depends on the dimensionality of the collective variables and the dimensionality (and type) of the orthogonal degrees of freedom, which is dramatically higher for the explicit water solution. For the dipeptide (Figure 6) MTD initially has lower error than ABF, although the difference disappears by 1 ns. The MTD/ICF/GPR method outperforms the other adaptive methods, by a large factor in the gas phase, and by a smaller margin in solution. To understand the role of equilibration, consider the BM and BM(only-sampling) curves. In the gas phase, including equilibration time (as in the BM curve) does not greatly increase the simulation time for a given accuracy. In solution, where there are many slow orthogonal degrees of freedom associated with the water molecules, the time needed for equilibration greatly increases the overall computational cost for a given accuracy.

For alanine tripeptide the same general ordering holds, but the advantage of MTD/ICF/GPR is much larger. It greatly improves the accuracy of the reconstructed free energy surface as compared to MTD and ABF: in both gas phase and solution a precision of ~ 1.0 kcal mol⁻¹ RMS over the entire 4 dimensional free energy surfaces is achieved with at least an order of magnitude shorter simulation time (Figure 7), getting much closer to the maximum obtainable accuracy of the artificial BM(only-sampling) limit. Combining the fast exploration of MTD with gradient sampling and robust reconstruction using GPR leads to a method that gives the free energy surface with higher accuracy with shorter simulation times.

For an alternate view of the reconstructed free energy surface simulation time dependence we show visualizations of several 3-D slices of the 4-D surfaces of the alanine tripeptide in solution, with the BM reference surface in Figure 8, and the reconstructed surfaces after 1, 5, and 10 ns in Figure 9. After 1 ns of simulation the ABF and MTD results still appear very different from the reference surface, while MTD/ICF/GPR already shows the correct general structure of minima and maxima. The ABF result improves systematically but slowly, at least partially due to the slow exploration shown in Figure 5. The MTD result, in contrast, becomes qualitatively reasonable before ABF, but remains noisy as expected for a histogram based method. Finally, we note that

detailed examination of the evolution of the local error (Figure 10) of the free energy surface in solution reflects the previously discussed characteristics of the methods: for metadynamics the fast exploration rate leads to a uniform error distribution within 5 ns, while for ABF the magnitude of the error strongly correlates with the free energy value. Since ABF has a more limited exploration rate regions having higher free energy values are visited less frequently, and the error in these regions dominates. Since in our MTD/ICF/GPR method the bias is governed by metadynamics it also provides a rather uniform error distribution after a few ns.

5 Conclusions

In this paper, we contrasted the two fundamentally different ways to reconstruct multidimensional free energy surfaces from sampled data: (i) using the location trajectory to estimate the equilibrium probability distribution in collective variable space, and (ii) collecting observations of an estimator for the free energy gradient and finding the corresponding potential of mean force. The two approaches form the basis for a wide range of free energy calculation methods, including metadynamics which uses (i) and adaptive biasing force which uses (ii). Comparing the gradient based and histogram based approaches in a Gaussian toy model with one collective variable revealed how the system parameters have very different effects on the error of reconstruction. The histogram-based reconstruction error depends only on the number of samples in each bin, and was therefore independent of the potential parameters. The gradient-based reconstruction error depends on the variance of the gradient estimator, which in turn depends on the potential parameters, and was smaller than the histogram-based reconstruction error for a wide range of parameters. Using gradient information to reconstruct the free energy surface in more than one dimension presents two challenges: noise due to the difficulty of obtaining enough samples, and the need for multidimensional integration. We use Gaussian process regression to address both of these, taking advantage of its ability to incorporate gradient samples directly without binning, and its robustness to noise.

In scientifically relevant systems the free energy surface is often needed in more than one or two dimensions, and one generally wants to avoid exploring all of collective variable space and rather focus on those regions with low free energy and low barriers that connect these. Therefore adaptive methods, which build up a biasing potential using partial free energy reconstructions from the past trajectory, are almost exclusively used. Tests on short peptides in gas phase and solution revealed that metadynamics has a significantly better rate of exploration than adaptive biasing force, owing to the repulsive hills that guide the trajectory away from its past locations. An approach that is superior to both MTD and ABF is to combine the biasing algorithm of metadynamics with the “instantaneous collective force” gradient estimator used in the adaptive biasing force method and reconstruct the final free energy surface using Gaussian process regression. The computational efficiency gain of this new MTD/ICF/GPR scheme is remarkable: it requires an order of magnitude shorter simulation time as compared with the other adaptive methods tested to reduce the RMS noise in the multidimensional free energy surface to a chemically accurate level.

Because the new scheme uses components of existing methods, it is possible to use existing software, e.g. the NAMD package. To carry out a simulation under this new protocol, both metadynamics and adaptive biasing force modules are turned on, but with the biasing part of the latter turned off. The instantaneous collective forces that are printed along the trajectory are then post-processed with GPR, for which a small software tool is available from https://github.com/molet/program_gp_general.

Acknowledgement

The work of N.B. was sponsored by Office of Naval Research through the Naval Research Laboratory’s 6.1 base program. G. C. acknowledges support from EPSRC under grant no. EP/J010847/1.

Supporting Information Available

Computational details of the harmonic oscillator models and alanine peptide systems, MD and GPR parameters of the reference surfaces, optimal MD and GPR parameters of the free energy surface methods (MTD and ABF), description of GPR sparsification and applied parameters, comparison of different implementation of MTD/ICF/GPR with varying thermostats and computational approximations of ICF.

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

References

- (1) Wang, F.; Landau, D. P. *Phys. Rev. Lett.* **2001**, *86*, 2050–2053.
- (2) Pártay, L. B.; Bartók, A. P.; Csányi, G. *The Journal of Physical Chemistry B* **2010**, *114*, 10502–10512.
- (3) Zwanzig, R. W. *The Journal of Chemical Physics* **1954**, *22*, 1420–1426.
- (4) Kirkwood, J. G. *The Journal of Chemical Physics* **1935**, *3*, 300–313.
- (5) Kirkwood, J. G. In *Theory of Liquids*; Alder, B. J., Ed.; Gordon and Breach: New York, 1968.
- (6) Kollman, P. *Chemical Reviews* **1993**, *93*, 2395–2417.
- (7) Mones, L.; Kulhánek, P.; Simon, I.; Laio, A.; Fuxreiter, M. *The Journal of Physical Chemistry B* **2009**, *113*, 7867–7873.
- (8) Rosta, E.; Woodcock, H. L.; Brooks, B. R.; Hummer, G. *Journal of Computational Chemistry* **2009**, *30*, 1634–1641.
- (9) Mones, L.; Csányi, G. *The Journal of Physical Chemistry B* **2012**, *116*, 14876–14885.
- (10) Rohrdanz, M. A.; Zheng, W.; Maggioni, M.; Clementi, C. *The Journal of Chemical Physics* **2011**, *134*, 124116.
- (11) Hashemian, B.; Millán, D.; Arroyo, M. *The Journal of Chemical Physics* **2013**, *139*, 214101.
- (12) Fiorin, G.; Klein, M. L.; Hnin, J. *Molecular Physics* **2013**, *111*, 3345–3362.
- (13) Cuendet, M. A.; Tuckerman, M. E. *Journal of Chemical Theory and Computation* **2014**, *10*, 2975–2986.
- (14) Laio, A.; Parrinello, M. *Proceedings of the National Academy of Sciences* **2002**, *99*, 12562–12566.

- (15) Laio, A.; Gervasio, F. L. *Reports on Progress in Physics* **2008**, *71*, 126601.
- (16) Darve, E.; Pohorille, A. *The Journal of Chemical Physics* **2001**, *115*, 9169–9183.
- (17) Stecher, T.; Bernstein, N.; Csányi, G. *Journal of Chemical Theory and Computation* **2014**, *10*, 4079–4097.
- (18) Lelièvre, T., Stoltz, G., Rousset, M., Eds. *Free Energy Computations: mathematical perspective*; Imperial College Press London, 2010.
- (19) Barducci, A.; Bussi, G.; Parrinello, M. *Phys. Rev. Lett.* **2008**, *100*, 020603.
- (20) Marsili, S.; Barducci, A.; Chelli, R.; Procacci, P.; Schettino, V. *The Journal of Physical Chemistry B* **2006**, *110*, 14011–14013.
- (21) Dickson, B. M.; Legoll, F.; Lelièvre, T.; Stoltz, G.; Fleurat-Lessard, P. *The Journal of Physical Chemistry B* **2010**, *114*, 5823–5830.
- (22) Hansen, H. S.; Hünenberger, P. H. *Journal of Computational Chemistry* **2010**, *31*, 1–23.
- (23) Huber, T.; Torda, A. E.; van Gunsteren, W. F. *Journal of Computer-Aided Molecular Design* **1994**, *8*, 695–708.
- (24) Grubmüller, H. *Phys. Rev. E* **1995**, *52*, 2893–2906.
- (25) Lelièvre, T.; Rousset, M.; Stoltz, G. *Nonlinearity* **2008**, *21*, 1155.
- (26) Alrachid, H.; Lelièvre, T. *ArXiv e-prints* **2015**, <http://arxiv.org/abs/1501.07094>.
- (27) Chipot, C., Pohorille, A., Eds. *Free Energy Calculations*; Springer-Verlag Berlin Heidelberg, 2007.
- (28) Ruiz-Montero, M. J.; Frenkel, D.; Brey, J. J. *Molecular Physics* **1997**, *90*, 925–942.
- (29) Darve, E.; Wilson, M. A.; Pohorille, A. *Molecular Simulation* **2002**, *28*, 113–144.


- (30) Rodriguez-Gomez, D.; Darve, E.; Pohorille, A. *The Journal of Chemical Physics* **2004**, *120*, 3563–3578.
- (31) Kulhánek, P.; Fuxreiter, M.; Koča, J.; Mones, L.; Střelcová, Z.; Petřek, M. <https://github.com/kulhanek/pmflib>, <https://lcc.ncbr.muni.cz/whitezone/development/pmflib>.
- (32) Hémin, J.; Chipot, C. *The Journal of Chemical Physics* **2004**, *121*, 2904–2914.
- (33) Hnin, J.; Fiorin, G.; Chipot, C.; Klein, M. L. *Journal of Chemical Theory and Computation* **2010**, *6*, 35–47.
- (34) Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kalé, L.; Schulten, K. *Journal of Computational Chemistry* **2005**, *26*, 1781–1802.
- (35) MacKay, D. J. *Information Theory, Inference and Learning Algorithms*; Cambridge University Press, Cambridge, UK, 2003.
- (36) Rasmussen, C. E.; Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series)*; MIT Press, Cambridge MA, 2005.
- (37) Quiñonero Candela, J.; Rasmussen, C. E. *J. Mach. Learn. Res.* **2005**, *6*, 1939–1959.
- (38) Mahoney, M. W.; Drineas, P. *Proceedings of the National Academy of Sciences* **2009**, *106*, 697–702.
- (39) Carter, E.; Ciccotti, G.; Hynes, J. T.; Kapral, R. *Chemical Physics Letters* **1989**, *156*, 472 – 477.
- (40) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. *Prot. Struct. Funct. Bioinf.* **2006**, *65*, 712–725.
- (41) MacKerell, A. D.; Bashford, D.; Bellott, M.; Dunbrack, R. L.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F.

- T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; Roux, B.; Schlenkrich, M.; Smith, J. C.; Stote, R.; Straub, J.; Watanabe, M.; Wirkiewicz-Kuczera, J.; Yin, D.; Karplus, M. *The Journal of Physical Chemistry B* **1998**, *102*, 3586–3616.
- (42) Darden, T.; York, D.; Pedersen, L. *The Journal of Chemical Physics* **1993**, *98*, 10089–10092.
- (43) Case, D. A.; Babin, V.; Berryman, J. T.; Betz, R. M.; Cai, Q.; Cerutti, D. S.; Cheatham, III, T. E.; Darden, T. A.; Duke, R. E.; Gohlke, H.; Goetz, A. W.; Gusarov, S.; Homeyer, N.; Janowski, P.; Kaus, J.; Kolossváry, I.; Kovalenko, A.; Lee, T. S.; LeGrand, S.; Luchko, T.; Luo, R.; Madej, B.; Merz, K. M.; Paesani, F.; Roe, D. R.; Roitberg, A.; Sagui, C.; Salomon-Ferrer, R.; Seabra, G.; Simmerling, C. L.; Smith, W.; Swails, J.; Walker, R. C.; Wang, J.; Wolf, R. M.; Wu, X.; Kollman, P. A. *AMBER 14*; 2014.


List of Figures

1	Alanine dipeptide (a) and tripeptide (b) and their dihedrals used in free energy surface calculations.	30
2	Two dimensional harmonic oscillator toy model joint probability density function of $x-y$ (first row), marginal distribution of x (second row) and conditional distribution of f_x at three different x values (third row, $x=0.0$ – orange, $x=1.0$ – green and $x=2.0$ – magenta) for three different ϕ values with $\sigma_1^2 = 1.0$ and $\sigma_2^2 = 0.04$ using one million samples.	31
3	Left panels: RMS errors of the histogram and ICF based free energy surface reconstructions (same colour map) for the 2 dimensional oscillator model as a function of potential parameters, evaluated for 32 bins. Right panel: ratio of the errors, with the black lines separating the regions where each method is superior.	32
4	Exploration rates of the entire free energy surface of alanine dipeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.	33
5	Exploration rates of the entire free energy surface of alanine tripeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.	34
6	RMS errors of the reconstructed free energy profiles of alanine dipeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.	35
7	RMS errors of the reconstructed free energy profiles of alanine tripeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.	36
8	3 dimensional slice of the reference free energy surface of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$, with the global minimum of the profile shifted to 0.	37

9	Reconstructed free energy surfaces of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$ as a function of simulation time. The global minimum of the profiles was shifted to 0 and the same colour map was applied as for the reference surface.	38
10	Time evolution of the local error of the reconstructed free energy surfaces of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$	39



(a)



(b)

Figure 1: Alanine dipeptide (a) and tripeptide (b) and their dihedrals used in free energy surface calculations.

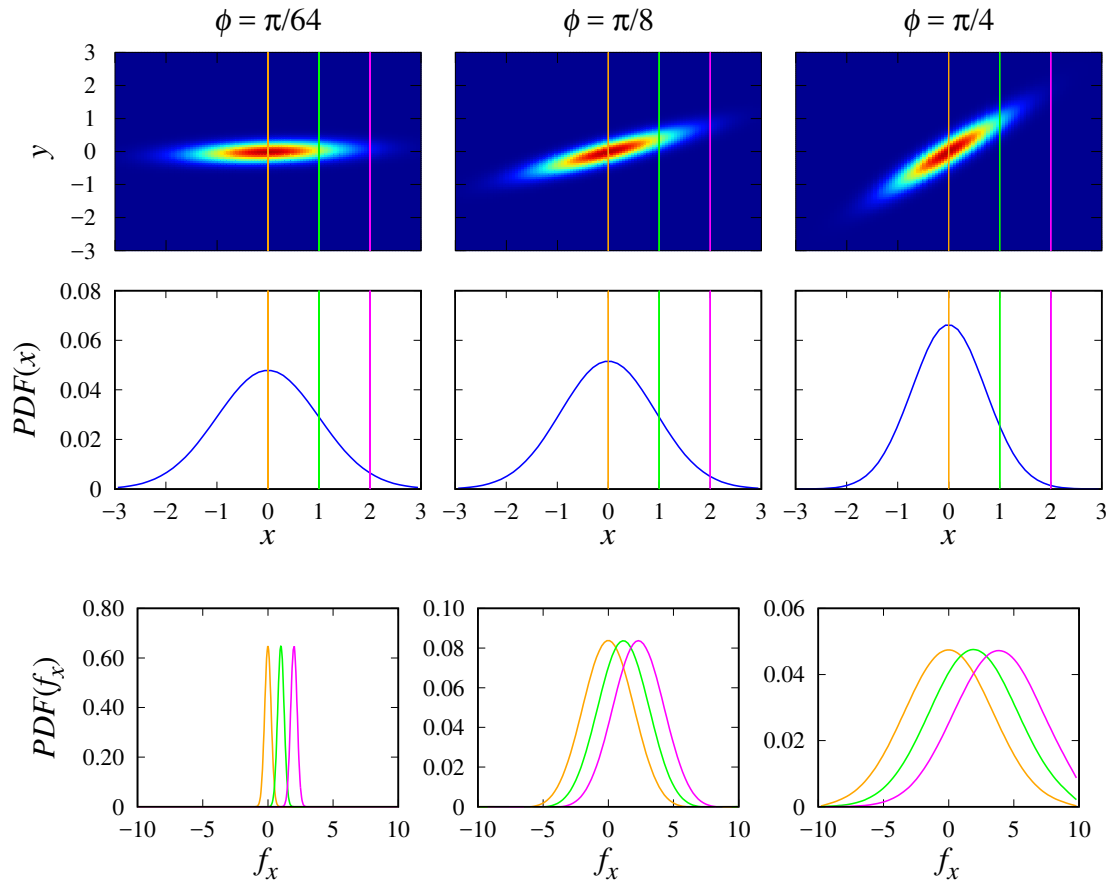


Figure 2: Two dimensional harmonic oscillator toy model joint probability density function of x - y (first row), marginal distribution of x (second row) and conditional distribution of f_x at three different x values (third row, $x=0.0$ – orange, $x=1.0$ – green and $x=2.0$ – magenta) for three different ϕ values with $\sigma_1^2 = 1.0$ and $\sigma_2^2 = 0.04$ using one million samples.




Figure 3: Left panels: RMS errors of the histogram and ICF based free energy surface reconstructions (same colour map) for the 2 dimensional oscillator model as a function of potential parameters, evaluated for 32 bins. Right panel: ratio of the errors, with the black lines separating the regions where each method is superior.

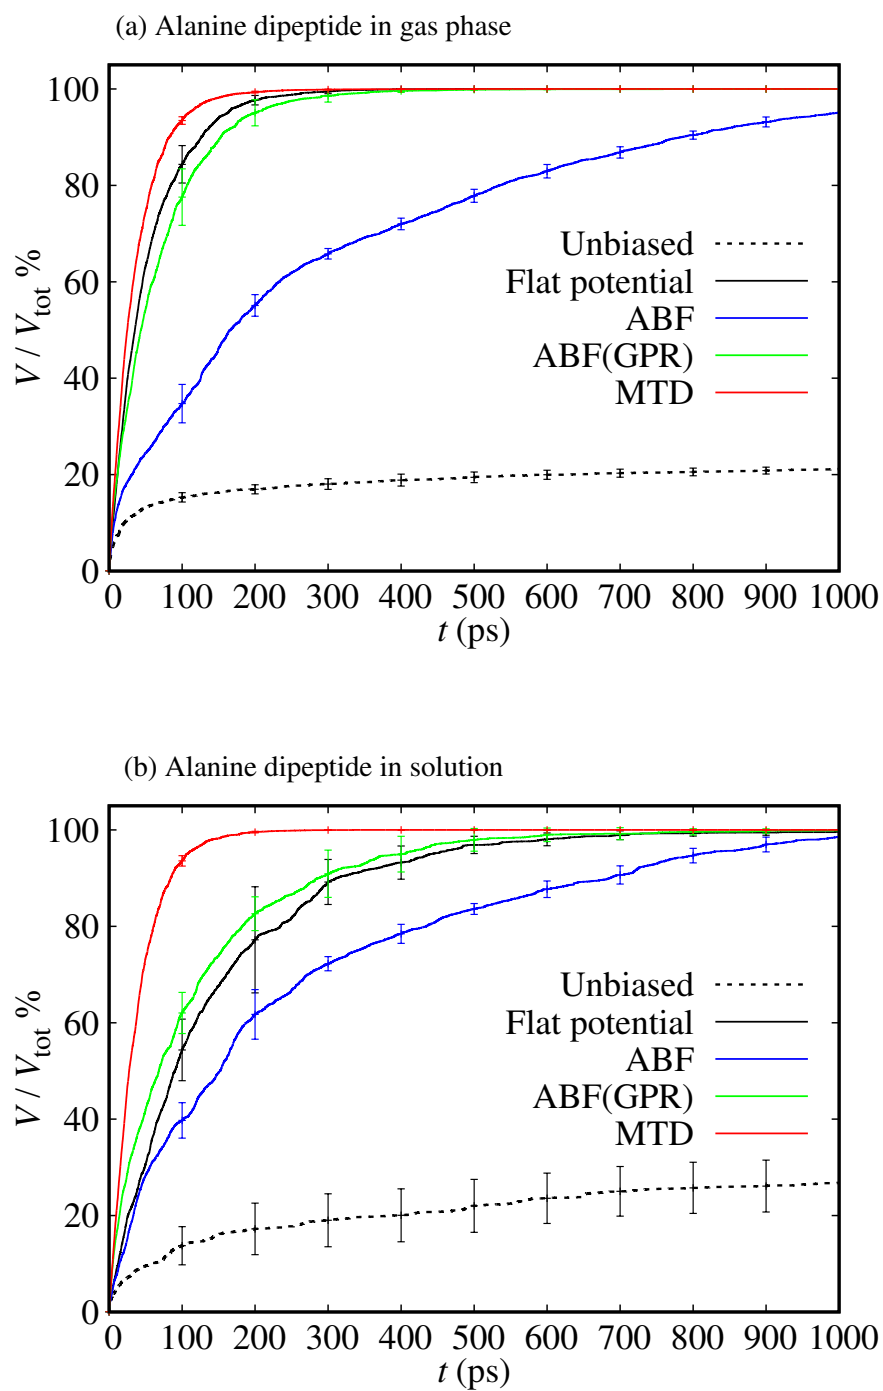


Figure 4: Exploration rates of the entire free energy surface of alanine dipeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.




Figure 5: Exploration rates of the entire free energy surface of alanine tripeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.

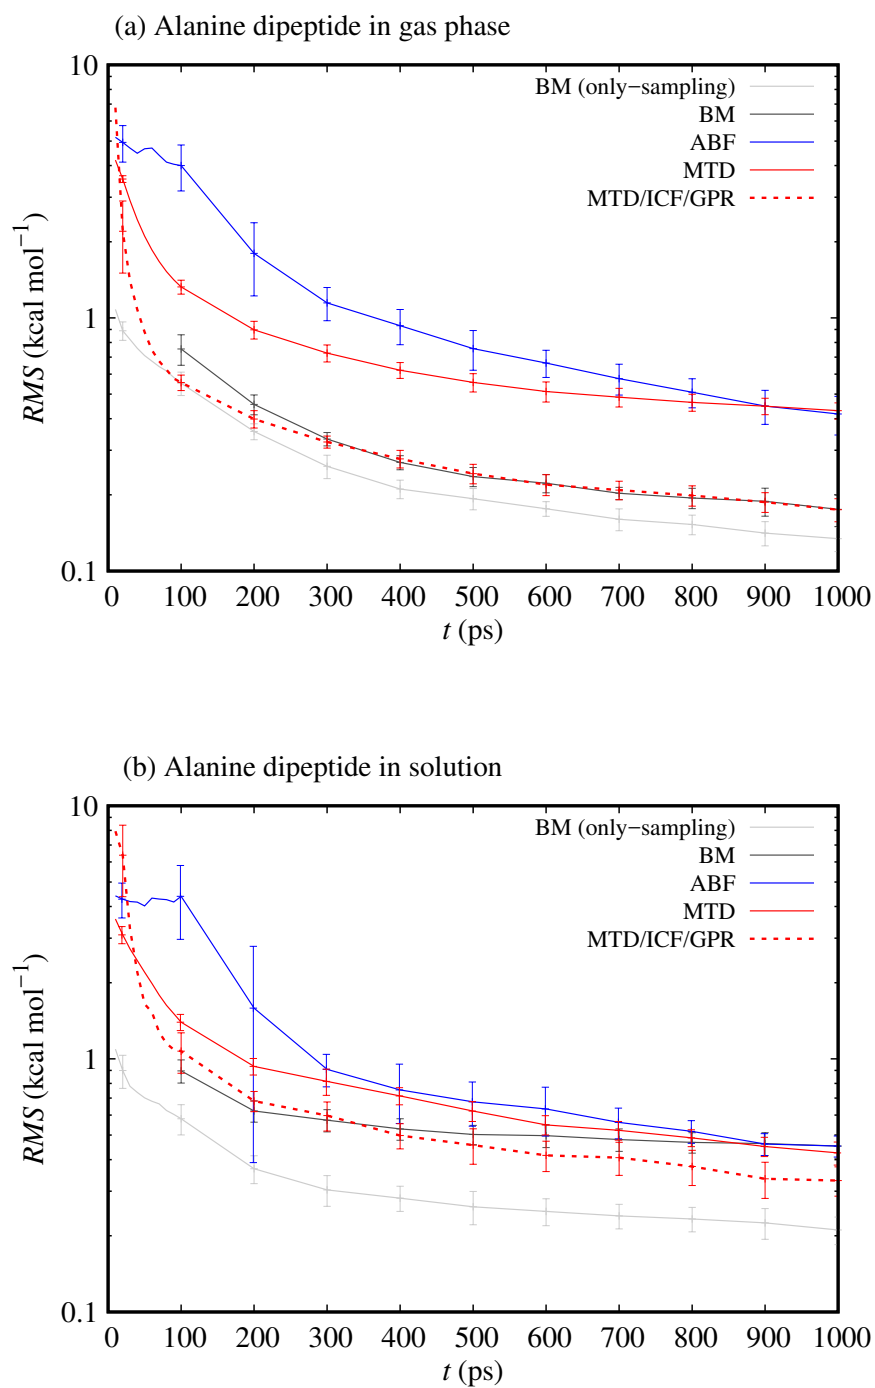


Figure 6: RMS errors of the reconstructed free energy profiles of alanine dipeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.

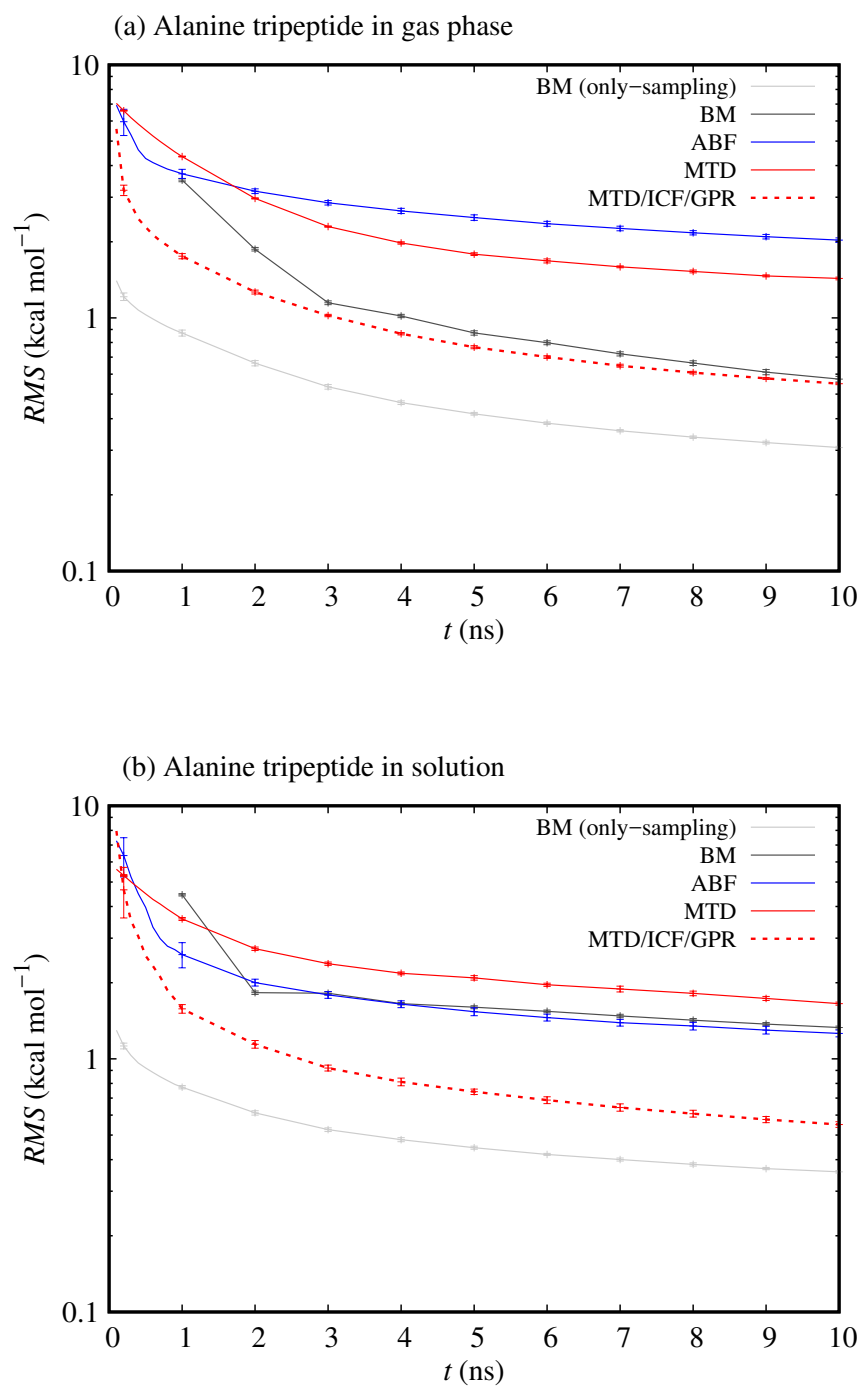


Figure 7: RMS errors of the reconstructed free energy profiles of alanine tripeptide in gas phase (a) and solution (b) using different methods. Error bars represent one standard deviation based on 10 independent simulations.




Figure 8: 3 dimensional slice of the reference free energy surface of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$, with the global minimum of the profile shifted to 0.

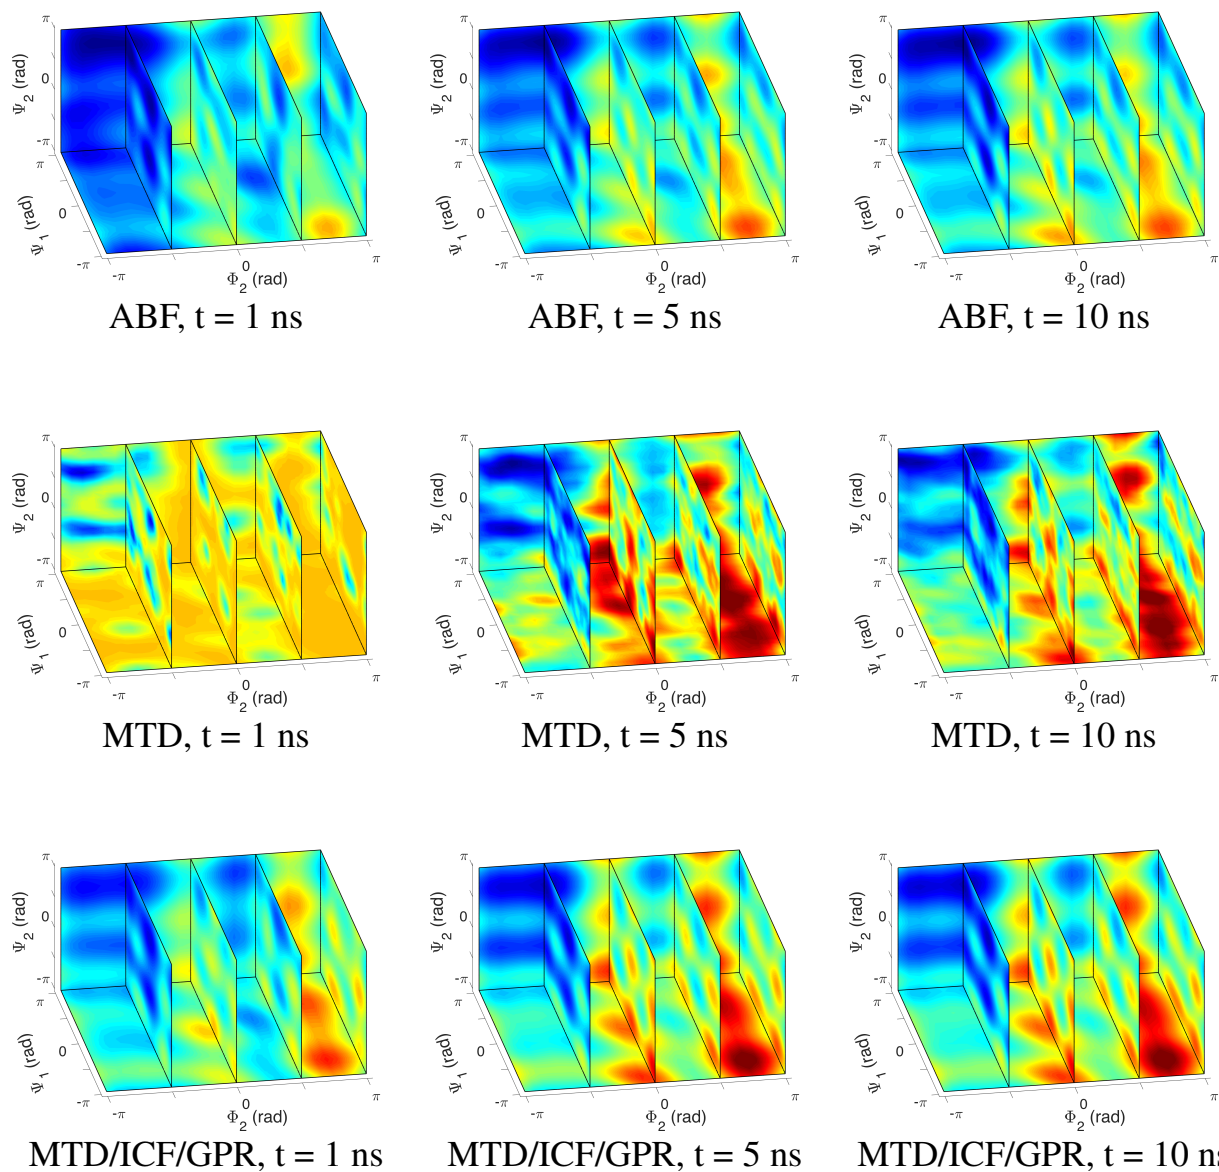


Figure 9: Reconstructed free energy surfaces of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$ as a function of simulation time. The global minimum of the profiles was shifted to 0 and the same colour map was applied as for the reference surface.

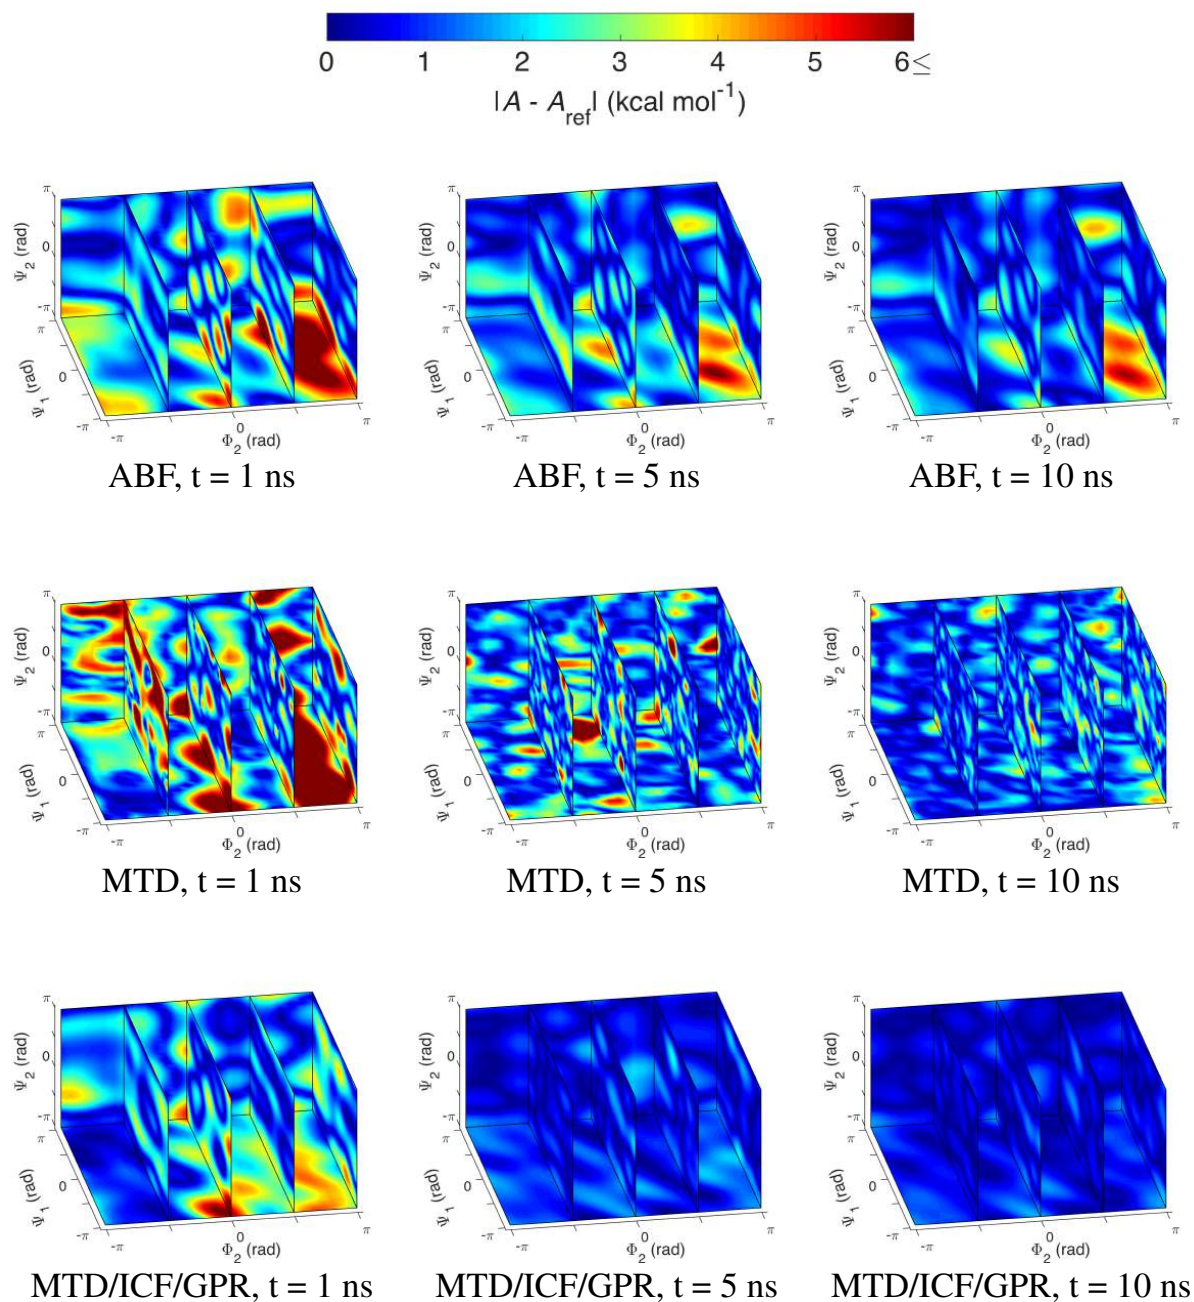


Figure 10: Time evolution of the local error of the reconstructed free energy surfaces of alanine tripeptide in solution at $\Phi_1 = 2\pi/25$.

Graphical TOC Entry

