Structural Bioinformatics

Bayesian active learning for optimization and uncertainty quantification in protein docking

Yue Cao and Yang Shen *

Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843, United States.

*To whom correspondence should be addressed.

Abstract

Motivation: *Ab initio* protein docking represents a major challenge for optimizing a noisy and costly “black box”-like function in a high-dimensional space. Despite progress in this field, there is no docking method available for rigorous uncertainty quantification (UQ) of its solution quality (e.g. interface RMSD or iRMSD).

Results: We introduce a novel algorithm, Bayesian Active Learning (BAL), for optimization and UQ of such black-box functions and flexible protein docking. BAL directly models the posterior distribution of the global optimum (or native structures for protein docking) with active sampling and posterior estimation iteratively feeding each other. Furthermore, we use complex normal modes to represent a homogeneous Euclidean conformation space suitable for high-dimension optimization and construct funnel-like energy models for encounter complexes. Over a protein docking benchmark set and a CAPRI set including homology docking, we establish that BAL significantly improve against both starting points by rigid docking and refinements by particle swarm optimization, providing for one third targets a top-3 near-native prediction. BAL also generates tight confidence intervals with half range around 25% of iRMSD and confidence level at 85%. Its estimated probability of a prediction being native or not achieves binary classification AUROC at 0.93 and AUPRC over 0.60 (compared to 0.14 by chance); and also found to help ranking predictions. To the best of knowledge, this study represents the first uncertainty quantification solution for protein docking, with theoretical rigor and comprehensive assessment.

Availability: Source codes are available at <https://github.com/Shen-Lab/BAL>.

Contact: yshen@tamu.edu

Supplementary information: https://github.com/Shen-Lab/BAL/tree/master/Paper_SI/

1 Introduction

Protein-protein interactions underlie many cellular processes, which has been increasingly revealed by the quickly advancing high-throughput experimental methods. However, compared to the binary information about what protein pairs interact, the structural knowledge about how proteins interact remains relatively scarce (Mosca *et al.*, 2013). Protein docking helps close such a gap by computationally predicting the 3D structures of protein-protein complexes given individual proteins’ 1D sequences or 3D structures (Smith and Sternberg, 2002).

Ab initio protein docking is often recast as an energy (or other objective functions) optimization problem. For the type of objective functions relevant to protein docking, neither the analytical form nor the gradient information would help global optimization as the functions are non-convex and extremely rugged thus their gradients are too local to inform global landscape. So the objective functions are often treated as *de facto* “black-box” functions for global optimization. Meanwhile, these functions

are very expensive to evaluate. Various methods, especially refinement-stage methods, have progressed to effectively sample the high-dimensional conformational space against the expensive functional evaluations (Gray *et al.*, 2003; Moal and Bates, 2010; Shen, 2013; Jiménez-García *et al.*, 2017; Marze *et al.*, 2018; Pfeifferberger and Bates, 2018).

When solving such optimization problems still remains a great challenge, quantifying the uncertainty of protein-docking results in some quality of interest (very often interface RMSD or iRMSD) is not addressed by any protein-docking method. Even though such information is much needed by the end users, current protein-docking methods often generate a rank-ordered list of results without giving quality estimation and uncertainty quantification to individual results and without providing the confidence in whether the entire list contains a quality result (for instance, a near-native protein-complex model with $iRMSD \leq 4\text{\AA}$).

Sources of uncertainty in protein docking methods include the objective function as well as the sampling scheme, which can be classified as epistemic uncertainty (Der Kiureghian and Ditlevsen, 2009). For instance, energy models as objective functions provide noisy and approximate observations of the assumed ground truth — the Gibbs

free energy; and iterative sampling techniques suffer from both the approximation of the search space (e.g. rotamerized side chains) and insufficient data in the approximated space (e.g. small numbers of samples considering the high dimensionality of the search space). In addition, uncertainty in protein structure data (e.g. X-ray crystal structures of proteins being "averaged" versions of their native conformations and derived from fitting observed diffraction patterns), which can be classified as aleatoric uncertainty, also enters protein docking methods when crystal structures are used as ground-truth native structures for training objective functions or tuning parameters in protein-docking methods.

Whereas the forward propagation of aleatoric uncertainty in protein structure data to structure-determined quantities has been studied empirically (Li *et al.*, 2016) and theoretically (Rasheed *et al.*, 2017), the much more difficult, inverse quantification of uncertainty in predicted protein or protein-complex structures originating from epistemic uncertainty in computational methods, is still lacking a mathematically rigorous solution. A unique challenge for uncertainty quantification (UQ) in protein docking is that the desired quality of interest here is directly determined by the optimum itself rather than the optimal value. In other words, closeness to native structures (for instance, measured by iRMSD) is an indicator for the usefulness of the docking results, but closeness to native structures' energy values is not necessarily the case. Therefore, UQ in protein docking has to be jointly solved with function optimization when finding the inverse mapping from a docking objective function to its global optimum is neither analytically plausible nor empirically cheap.

In this study, we introduce a rigorous Bayesian framework to simultaneously perform function optimization and uncertainty quantification for expensive-to-evaluate black-box objective functions. To that end, our Bayesian active learning (BAL) iteratively and adaptively generate samples and update posterior distributions of the global optimum. We propose a posterior in the form of the Boltzmann distribution with adaptive annealing schedule that gradually shifts the search from exploration to exploitation and with an objective-function estimator (surrogate) that is a non-parametric Kriging regressor. The iteratively updated posterior carries the belief (and uncertainty as well) on where the global optimum is given historic samples and guides next-iteration samples, which presents an efficient data-collection scheme for both optimization and UQ. Compared to typical Bayesian optimization methods (Shahriari *et al.*, 2016) that first model the posterior of the objective function and then optimize the resulting functional, our BAL framework directly models the posterior for the global optimum and overcomes the intensive computations in both steps of typical Bayesian optimization methods. Compared to another work (Ortega *et al.*, 2012) that also models the posterior of the global optimum, we provide both theoretical and empirical results that our BAL has a consistent and unbiased estimator as well as a global uncertainty-aware and dimension-dependent annealing schedule.

We also make innovative contributions in the application domain of protein docking. Specifically, we design a machine learning-based objective function that estimates binding affinities for docked encounter complexes as well as assesses the quality of interest, iRMSD, for docking results. We also re-parameterize the search space for both external rigid-body motions (Shen *et al.*, 2008) and internal flexibility, into a low-dimensional homogeneous and isotropic space suitable for high-dimensional optimization, using our (protein) complex normal modes (cNMA) (Oliwa and Shen, 2015). Considering that protein docking refinement often starts with initial predictions representing separate conformational clusters/regions, we use estimated local posteriors over individual regions to construct local and global partition functions; and then calculate the probability that the prediction for each conformational cluster, each conformational cluster, or the entire list of conformational clusters is near-native.

The rest of the paper is organized as following. In Materials and Methods, we first give a mathematical formulation for the optimization and the UQ, then introduce our Bayesian active learning (BAL) that iteratively update sampling and posterior estimation. We next introduce the parameterization of the search space that allows concurrent and homogeneous sampling of external rigid-body and internal flexible-body motions as well as newly-developed machine learning models as the noisy energy function that estimates the binding free energy for encounter complexes. And we end the Materials and Methods with uncertainty quantification for protein docking. In Results and Discussion, using a comprehensive protein docking benchmark set involving unbound docking and a CAPRI set involving homology docking, we assess optimization results for BAL with comparison to starting structures from ZDOCK and refined structures by particle swarm optimization (PSO). We further assess the uncertainty quantification results: accuracy of the confidence levels and tightness of the confidence region, as well as the confidence scores on the near-nativeness of predictions. Lastly, before reaching Conclusions, we visualize the estimated energy landscape and confirm that the funnel-like energy landscapes do exist near native structures in the homogeneous conformational space blending external rigid-body and internal flexible-body motions.

2 Materials and Methods

2.1 Mathematical Formulation

We consider a black-box function $f(\mathbf{x})$ (e.g. ΔG , the change in the Gibbs free energy upon protein-protein interaction) that can only be evaluated at any sample with an expensive yet noisy observation $y(\mathbf{x})$ (e.g. modeled energy difference). Our goal in optimization is

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

(e.g. the native structure of a protein complex). And our goal in uncertainty quantification is the probability that $\hat{\mathbf{x}}$, a prediction of $\mathbf{x}^*$, falls in an interval $[lb, ub]$ of quality relative to $\mathbf{x}^*$:

$$P(lb \leq Q(\hat{\mathbf{x}}, \mathbf{x}^*) \leq ub) = 1 - \sigma$$

where $1 - \sigma$ is the confidence level; $[lb, ub]$ is the confidence interval in the solution quality; and $Q(\cdot, \cdot)$, the quality of interest measuring some distance or dissimilarity, can be an Euclidean norm (as in our assessment for test functions), another distance metric, or other choices of users (for instance, iRMSD as in our assessment for protein docking with $ub = 4\text{\AA}$).

2.2 Bayesian active learning with a posterior of $\mathbf{x}^*$

We address the problem above in a Bayesian perspective: instead of treating $\mathbf{x}^*$ as a fixed point, we model $\mathbf{x}^*$ as a random variable and construct its probability distribution, $p(\mathbf{x}^*|\mathcal{D})$, given samples $\mathcal{D} = \{(\mathbf{x}, y)\}$. This probability distribution, carrying the belief and the uncertainty on the location of $\mathbf{x}^*$, is a prior when $\mathcal{D} = \emptyset$ (no sample) and a posterior otherwise. Considering the cost of function evaluation, we iteratively collect new samples in iteration t (where all samples collected by the end of the t -th iteration are denoted $\mathcal{D}^{(t)}$) based on the latest estimated posterior, $p(\mathbf{x}^*|\mathcal{D}^{(t-1)})$; and we update the posterior $p(\mathbf{x}^*|\mathcal{D}^{(t)})$ based on $\mathcal{D}^{(t)}$. An illustration of the iterative approach is given in Fig. 1.

For optimization, we set $\hat{\mathbf{x}}$ to be the best sample with the lowest y value given a computational budget (reflected in the number of samples or iterations). For UQ, given the posterior $p(\mathbf{x}^*|\mathcal{D}^{(t)})$ in the final iteration, one can propagate the inferred uncertainty in $\mathbf{x}^*$ forwardly to that in the quality of interest, $Q(\hat{\mathbf{x}}, \mathbf{x}^*)$, for the found $\hat{\mathbf{x}}$, using techniques such as Markov chain Monte Carlo.

Fig. 1: Illustration of the Bayesian active learning (BAL) algorithm. (A): A typical energy landscape projected onto the first principle component (PC1). The dashed line indicates the location of the optimal solution. (B)-(D): The samples (dots) and the kriging regressors (light curves) at the stage of the 1st, 4th and 10th iteration, respectively. Samples are colored from cold to hot by the iterations and those in the same iteration have the same color. (E): The entropy (measuring uncertainty) of the posterior reduces as the number of samples increases. Its quick drop from 30 to 100 samples corresponds to a drastic change of kriging regressor, which suggests increasing exploitation in possible function basins. After 100 samples, the entropy goes down more slowly, meaning that the regressor only has a small amount of adjustment. (F)-(H): The corresponding posterior distributions for (B)-(D), respectively.

2.2.1 Non-parametric posterior of $\mathbf{x}^*$

We propose to use the Boltzmann distribution to describe the posterior

$$p(\mathbf{x}^* | D^{(t)}) \propto \exp(-\rho \cdot \hat{f}(\mathbf{x}))$$

where $\hat{f}(\mathbf{x})$ is an estimator for $f(\mathbf{x})$, and ρ is a parameter (sometimes $1/RT$ where R is the gas constant and T the temperature of the molecular system).

To iteratively guide the expensive sampling and balance between exploration and exploitation in a data efficient way, we choose ρ to follow an adaptive annealing schedule over iteration t :

$$\rho_t = \rho_0 \cdot \exp((h_p^{(t-1)})^{-1} n_t^{-\frac{1}{d}})$$

where ρ_0 , the initial ρ , is a parameter; $h_p^{(t-1)}$ is the (continuous) entropy of the last-iteration posterior, a shorthand notation for $h(p(\mathbf{x}^* | D^{(t-1)}))$; $n_t = |D^{(t)}|$ is the number of samples collected so far; and d is the dimensionality of the search space $\mathcal{X}$.

This annealing schedule is inspired by the adaptive simulated annealing (ASA) (Ingber, 2000), especially the exponential form and the $n_t^{-\frac{1}{d}}$ term. However, we use the $(h_p^{(t-1)})^{-1}$ term rather than a constant as in ASA so that we exploit all historic samples $D^{(t)}$. In this way, as the uncertainty of $\mathbf{x}^*$ decreases, ρ_t increases and shifts the search toward exploitation.

The function estimator $\hat{f}(\mathbf{x})$ also updates iteratively according to the incrementally increasing n_t samples $D^{(t)} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n_t}$. We use a consistent and unbiased Kriging regressor (Chilès and Delfiner, 2012) which is known to be the best unbiased linear estimator (BLUE):

$$\hat{f}(\mathbf{x}) = f_0(\mathbf{x}) + (\boldsymbol{\kappa}(\mathbf{x}))^T (\mathbf{K} + \epsilon^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{f}_0)$$

where $f_0(\mathbf{x})$ is the prior for $E[f(\mathbf{x})]$; $\boldsymbol{\kappa}(\mathbf{x})$ is the kernel vector with the i th element being the kernel, a measure of similarity, between $\mathbf{x}$ and $\mathbf{x}_i$; $\mathbf{K}$ is the kernel matrix with the (i, j) element being the kernel between $\mathbf{x}_i$ and $\mathbf{x}_j$; $\mathbf{y}$ and $\mathbf{f}_0$ are the vector of $y_1, \dots, y_{n_t}$ and $f_0(\mathbf{x}_1), \dots, f_0(\mathbf{x}_{n_t})$, respectively; and ϵ reflects the noise in the observation and is estimated to be 2.1 as the prediction error for the training set.

We derive the Kriging regressor in the Supporting Information (SI) Sec. 1.2.2. And we will use the regressor to evaluate binding energy and estimate iRMSD for UQ over multiple regions in Sec. 2.5.

2.2.2 Adaptive sampling based on the latest posterior

For a sequential sampling policy that balances exploration and exploitation during the search for the optimum, we choose Thompson sampling (Russo *et al.*, 2017) which samples a batch of points in the t -th iteration based on the latest posterior $p(\mathbf{x}^* | D_{t-1})$. This seemingly simple policy has been found to be theoretically (Agrawal and Goyal, 2011) and empirically (Chapelle and Li, 2011) competitive compared to other updating policies such as Upper Confidence Bound (Shahriari *et al.*, 2016). In our case, it is actually straightforward to implement given the posterior on $\mathbf{x}^*$.

There are multiple reasons to collect in each iteration a batch of samples rather than a single one. First, given the high dimension of the search space, it is desired to collect adequate data before updating the posterior. Second, the batch sampling weakens the correlation among samples and make them more independent, which benefits the convergence rate of the Kriging regressor. Last, parallel computing could be trivially applied for batch sampling, which would significantly improve the algorithm throughput.

Fig. 1 gives an illustration of the algorithm behavior. The initial samples drawn from a uniform distribution leads to a relatively flat posterior whose maximum is off the function optimum (Fig. 1F). As the iteration progresses, the uncertainty about the optimum gradually reduces (Fig. 1E) and newer samples are increasingly focused (Fig. 1C,D) as the posteriors are becoming narrower with peaks shifting toward the function optimum (Fig. 1G,H).

In our docking study, $d = 12$ for a homogeneous space spanned by complex normal modes (see Sec. 2.3). We construct a prior and collect 30 samples in the first iteration and 20 in each of the subsequent iterations. We limit the number of iterations (samples) to be 31 (630) as a way to impose a computational budget.

2.2.3 Kernel with customized distance metric

The kernel in the Kriging regressor for the posterior is a measure of similarity. For test functions defined in an Euclidean space, we use the radial basis function (RBF) kernel:

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2l^2}\right)$$

where $\|\mathbf{x}_1 - \mathbf{x}_2\|^2$, a measure of dissimilarity, is the Euclidean distance and l , the bandwidth of the kernel, is set as $l = l_0 \cdot n_t^{\frac{1}{d}}$ following (Györfi *et al.*, 2002). l_0 , dependent on search space, is set at 2.0 for docking without particular optimization.

For protein docking, we replace the Euclidean distance in the RBF kernel with the interface RMSD (iRMSD) between two sample structures. iRMSD captures sample dissimilarity relevant to function-value dissimilarity and is independent of search-space parameterization.

For this purpose, we also have to address two technical issues. First, protein interface information is determined by $\mathbf{x}^*$ and thus unknown. We instead use the putative interface seen in the samples (see more details in the SI Sec. 2.4). Second, kernel calculation with iRMSD is time consuming. The time complexity of iRMSD calculation is $O(N)$ and that of regressor update is $O(Nn^2)$, where N , the number of interfacial atoms, can easily reach hundreds or thousands, and n , the number of samples, can also be large. To save computing time, we develop a fast RMSD calculation method that reduces its time complexity from $O(N)$ down to $O(1)$ (see details in SI Sec. 2.1).

2.2.4 Related methods

Current Bayesian optimization methods typically model the posterior distribution of $f(\mathbf{x})$ rather than that of $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ directly. After modeling the posterior distribution over the functional space (a common non-parametric way is through Gaussian processes), they would subsequently sample the functional space and optimize sample functions. For instance, Villemonteix *et al.* (2006) used Monte Carlo sampling; and Henrández-Lobato *et al.* (2014) discretized the functional space to approximate the sample paths of Gaussian processes using a finite number of basis functions then optimized each sample path to get one sample of $\mathbf{x}^*$. The two-step approach of current Bayesian optimization methods involve intensive sampling and (non-convex) optimization that is computationally intensive and not pragmatically useful for protein docking.

To our knowledge, Ortega *et al.* presented the only other study that directly models the posterior distribution over the optimum (Ortega *et al.*, 2012). Both methods fall in the general category of Bayesian optimization and use consistent non-parametric regressors. However, we prove in Sec. 1.2 of the SI that their regressor is biased whereas our Kriging regressor is unbiased. We explain in Sec. 1.1 there that their annealing schedule (temperature control) only considers the pairwise distance between samples without location awareness and is independent of dimensionality d ; whereas ours has a term involving location-aware global uncertainty and generalizes well to various dimensions. Beyond those theoretical comparison, we also included empirical results to show the superior optimization and UQ performances for BAL.

The rest of the Materials and Methods section involve methods specific to the protein docking problem: parameterization, dimensionality reduction, and range reduction of the search space $\mathcal{X}$; machine learning model as $y(\mathbf{x})$, i.e., an energy model for encounter complexes; uncertainty quantification for a predicted structure or a list of predictions as well as the corresponding quality estimation for ranking predictions or classifying their nativeness.

2.3 Conformational Sampling in $\mathcal{X}$

In protein docking the full search space $\mathcal{X}$ captures the degrees of freedom for all atoms involved. Let one protein be receptor whose position is fixed and the other be ligand. And let N_R , N_L , and N be the number of atoms for the receptor, the ligand, and the complex respectively. Then $\mathcal{X} = \mathbb{R}^{3N-6}$ is a Euclidean space whose dimension easily reaches 10^4 for a small protein complex without surrounding solvent molecules.

If accuracy is sacrificed for speed, proteins can be (unrealistically) considered rigid and $\mathcal{X} = SE(3) = \mathbb{R}^3 \times SO(3)$ (ligand translations and rotations) is a Riemannian manifold. Docking methods fall in the spectrum

between these two ends that are represented by all-atom molecular dynamics and FFT rigid docking, respectively. For instance, one can consider locally rigid pieces of a protein rather than a globally rigid protein, then $\mathcal{X}$ becomes the product of many $SE(3)$ for local rigidity (Mirzaei *et al.*, 2015); or one can model individual proteins' internal flexible-body motions normal modes on top of the ligand rigid-body motions, thus $\mathcal{X}$ becomes the product of $\mathbb{R}^K$ (where $K \ll N_{RL}$) and $SE(3)$ (Moal and Bates, 2010).

From the perspective of optimization and UQ, both the high-dimensionality of $\mathbb{R}^{3N-6}$ and the geometry of the lower-dimensional manifold present challenges. Almost all dimensionality reduction efforts in protein docking impose conditions (such as aforementioned local or global rigidity) in the full Euclidean space and lead to embedded manifolds difficult to (globally) optimize over. The challenge from the manifold has been either disregarded in protein docking or addressed by the local tangent space (Shen *et al.*, 2007, 2008; Mirzaei *et al.*, 2015).

Could and how could the dimensionality of the conformational space be reduced while its geometry maintains homogeneity and isotropy of a Euclidean space and its basis vectors span conformational changes of proteins upon interactions? In this subsection we give a novel approach to answer this question for the first time. In contrast to common conformational sampling that separates internal flexible-body motions (often Euclidean) and external rigid-body motions (a manifold) (Marze *et al.*, 2018), we re-parameterize the space into a Euclidean space spanned by complex normal modes Oliwa and Shen (2015) blending both flexible- and rigid-body motions. The mapping preserves distance metric in the original full space. We further reduce the dimensionality and the range in the resulting space.

2.3.1 Complex normal modes blend flexible- and rigid-body motions

We previously introduced complex normal mode analysis, cNMA (Oliwa and Shen, 2015), to model conformational changes of proteins during interactions. Using encounter complexes from rigid docking, cNMA extends anisotropic network model (ANM) to capture both conformational selection and induced fit effects. After the Hessian matrix is projected to remove the rigid-body motion of the receptor, its non-trivial eigenvectors $\boldsymbol{\mu}_j$ ($j = 1, \dots, 3N - 6$) form orthonormal basis vectors. We showed that $\boldsymbol{\mu}_j^R$, the components of the complex normal modes, better capture the direction of individual proteins' conformational changes than conventional NMA did (Oliwa and Shen, 2015). We also showed that the re-scaled eigenvalues for these components, $\lambda_j^R = \frac{\lambda_j}{\|\boldsymbol{\mu}_j^R\|^2}$, can be used to construct features for machine learning and predict the extent of the conformational changes.

2.3.2 Dimensionality reduction

In this study we focus on the motions of a whole complex rather than individual proteins and develop sampling techniques for protein docking. Each single complex normal mode simultaneously captures concerted flexible-body motions of individual proteins (receptor and ligand) and rigid-body motion of the non-fixed protein (ligand). They together span a homogeneous and isotropic Euclidean space where the distance between two points is exactly the RMSD between corresponding complex structures.

For dimensionality reduction in the resulting space, we choose the first K_1 non-trivial eigenvectors $\boldsymbol{\mu}_j$ ranked by increasing eigenvalues λ_j ; and we additionally include K_2 $\boldsymbol{\mu}_j$ (not in the first K_1) ranked by increasing λ_j^R . In other words, we choose K_1 slowest modes for the complex and another K_2 for the receptor as the set of basis vectors $\mathcal{B}$. K_1 and K_2 are set at 9 and 3, respectively, leading to the dimension of the reduced space to $d = 12$. A supplemental video illustrates the motion of the slowest such modes.

2.3.3 Range reduction

For range reduction in the dimension-reduced space, we perturb a starting complex structure $\vec{C} \in \mathbb{R}^{3N-6}$ along aforementioned basis vectors to generate sample $\vec{C}_0$ while enforcing a prior on the scaling factor s in the first iteration. Specifically

$$\vec{C} = \vec{C}_0 + \sum_{j \in \mathcal{B}} r_j \frac{s}{\sqrt{\lambda_j}} \cdot \boldsymbol{\mu}_j$$

where r_j , the coefficient of the j th normal mode $\boldsymbol{\mu}_j$, is uniformly sampled on S^d , the surface of a d -dimensional standard sphere with a unit radius. The scaling factor s is given by

$$s = \frac{\tau_R}{\frac{1}{\sqrt{N_R}} \left\| \sum_{j \in \mathcal{B}} \frac{r_j}{\sqrt{\lambda_j}} \cdot \boldsymbol{\mu}_j^R \right\|}$$

where τ_R is the estimated conformational change (measured by RMSD) between the unbound and the bound receptor. Note that $\boldsymbol{\mu}_i^R$'s are not orthonormal to each other.

We previously predicted τ_R by a machine learning model giving $\widehat{\text{RMSD}}_R$, a single value for each receptor (Chen *et al.*, 2017). Here we replace $\widehat{\text{RMSD}}_R$ with a predicted distribution by multiplying it to a truncated normal distribution $N(\mu = 0.99, \sigma^2 = 0.096)$ within $[0, 2.5]$. The latter distribution is derived by fitting the ratios between the actual and the predicted values, $\text{RMSD}_R / \widehat{\text{RMSD}}_R$, for 50 training protein complexes (see more details about datasets in Sec. 2.7 and those about distribution fitting in Sec. 2.2 of the **SI**). Therefore, our parameterization produces $\mathbf{x} = s \cdot \mathbf{r} \in \mathbb{R}^d$ whose prior is derived as above.

Since the ligand component of complex normal modes include simultaneous flexible- and rigid-body motions, conformational sampling could lead to severely distorted ligand geometry. We thus further restrict the ligand perturbation Δ_L (flexible- and rigid-body together) to be within $\overline{\Delta}_L$

$$\Delta_L = \sqrt{\frac{1}{N_L}} \left\| \sum_{j \in \mathcal{B}} r_j \frac{s}{\sqrt{\lambda_j}} \cdot \boldsymbol{\mu}_j^L \right\| \leq \overline{\Delta}_L$$

We set $\overline{\Delta}_L$ at 6Å according to the average size of binding energy attraction basins seen in conformational clusters (Kozakov *et al.*, 2005). For samples generated from the aforementioned prior or the updated posterior, we reject those violating the ligand perturbation limit. We discuss about the feasibility of the search region in **SI** Sec. 2.3.

2.4 Energy Model $y(\mathbf{x})$ for Sampling

2.4.1 Binding affinity prediction for sampled encounter complexes

We introduce a new energy model based on binding affinities $K'_d(\mathbf{x})$ of structure samples $\mathbf{x}$ that are often encounter complexes. The model assumes that K'_d correlates with K_d , the binding affinity of the native complex, and deteriorates with the increase of the sample's iRMSD (the encounter complex is less native-like):

$$K'_d(\mathbf{x}) = K_d \cdot \exp(\alpha \cdot (\text{iRMSD}(\mathbf{x}))^q)$$

where α and q are hyper-parameters optimized through cross-validation. In other words, we assume that the fraction of binding affinity loss is exponential in a polynomial of iRMSD. Therefore, the binding energy, a machine learning model $y(\mathbf{x}; \mathbf{w})$ of parameters $\mathbf{w}$ can be represented as

$$y(\mathbf{x}; \mathbf{w}) = RT \ln(K'_d(\mathbf{x})) = RT \ln(K_d) + RT\alpha \cdot (\text{iRMSD}(\mathbf{x}))^q$$

Given an observed or regressed $y(\mathbf{x})$ value, one can estimate $\text{iRMSD}(\mathbf{x})$ with given K_d .

2.4.2 Machine learning

We train machine learning models, including ridge regression with linear or RBF kernel and random forest, for $y(\mathbf{x}; \mathbf{w})$. The 8 features include changes upon protein interaction in energy terms such as internal energies in bond, angle, dihedral and Urey-Bradley terms, van der Waals, non-polar contribution of the solvation energy based on solvent-accessible surface area (SASA), and electrostatics modeled by Generalized Born with a simple SWitching (GBSW), all of which are calculated in a CHARMM27 force field.

We use the same training set of 50 protein pairs (see details in **SI** Sec. 2.5) as in predicting the extent of conformational change. From rigid docking and conformational sampling we generate 13,004 complex samples for 50 protein pairs, including 6,464 near-native and 6,540 non-native examples in the training set. Hyper-parameters of ridge regression with RBF kernel as well as random forest are optimized by cross-validation. And model parameters $\mathbf{w}$ are trained again over the entire training set with the best hyper-parameters. More details can be found in Sec. 2.6 of the **SI**. For the assessment, we use the test set **a** of 26 protein pairs (again in **SI** Sec. 2.5) and generate 20 samples similarly for each of the 10 initial docking results for each pair, leading to 5,200 cases.

2.5 Uncertainty Quantification for Protein Docking

A unique challenge to protein docking refinement is that, instead of optimization and UQ in a single region $\mathcal{X}$, we may do so in K separate ones $\mathcal{X}_i$ ($i = 1, \dots, K$) where each $\mathcal{X}_i$ is a promising conformational region/cluster represented by a initial docking result. This is often necessitated by the fact that the extremely rugged energy landscape is populated with low-energy basins separated by frequent high-energy peaks in a high-dimensional space, thus preferably searched over multiple stages (Vajda and Kozakov, 2009).

Our goal of UQ for protein docking results is to determine, for each $\hat{\mathbf{x}}_i$ – the prediction in $\mathcal{X}_i$ (the i th structure model), its quality bounds $[lb, ub]$ such that

$$P(lb \leq Q(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq ub) = 1 - \sigma$$

where the quality of interest $Q(\hat{\mathbf{x}}, \mathbf{x}^*)$ here is iRMSD between a predicted and the native structure and $1 - \sigma$ a desired confidence level.

To that end, we forwardly propagate the uncertainty from $\mathbf{x}^*$ (native structure) to iRMSD, given the final posterior $p(\mathbf{x}^* | \mathcal{D}_t)$ in individual regions (local posteriors). Specifically, we generate 1,000,000 samples following the local posterior using Markov chain Monte Carlo, evaluate their binding energies using the Kriging regressor, and estimate their iRMSD using our binding affinity prediction formula inversely. We then use these sample iRMSD values to determine confidence intervals $[lb, ub]$ for various confidence score $1 - \sigma$ so that $P(\text{iRMSD} < lb) = P(\text{iRMSD} > ub) = \sigma/2$.

2.6 Confidence scores for near-nativeness

We next calculate the probability that a prediction $\hat{\mathbf{x}}_i$ is near-native, i.e., $P(Q(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4)$ (Méndez *et al.*, 2005). Calculating this quantity would demand the probability that the native structure lies in the i th conformational region / cluster, $P(\mathbf{x}^* \in \mathcal{X}_i)$ ($P(\mathcal{X}_i)$ in short) as well as that the probability that it lies in all the K regions, $P(\mathbf{x}^* \in \cup_{i=1}^K \mathcal{X}_i)$ ($P(U_K)$ in short). By following the chain rule we easily reach

$$\begin{aligned} P(\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4) &= \\ P(\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4 | \mathcal{X}_i) \cdot P(\mathcal{X}_i | U_K) \cdot P(U_K) \end{aligned}$$

. Here we use the fact that $\{\mathbf{x}^* \in \mathcal{X}_i\} \subset \{\mathbf{x}^* \in \cup_{i=1}^K \mathcal{X}_i\}$ and assume that $\{\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4\} \subseteq \{\mathbf{x}^* \in \mathcal{X}_i\}$ (the range of conformational clusters in iRMSD is usually wider than 4Å).

We discuss how to calculate each of the three terms for the product.

2.6.1 $P(\text{iRMSD}(\hat{x}_i, x^*) \leq 4 | \mathcal{X}_i)$

If the native structure x^* (unknown) is contained in the i th region/cluster $\mathcal{X}_i$, what is the chance that the predicted structure $\hat{x}_i$ (known) is within 4Å? We again use forward uncertainty propagation starting with the local posterior $p(x^* | \mathcal{D}_i)$ in $\mathcal{X}_i$. We sample 100,000 structures following the posterior with Markov chain Monte Carlo, calculate their iRMSD to the prediction $\hat{x}_i$, and empirically determine the portion within 4Å for the probability of interest here. Notice that the native interface is unknown thus the putative interface is used instead.

2.6.2 $P(\mathcal{X}_i | U_K)$

If the native structure is contained in at least one of the K regions, what is the chance that it is in $\mathcal{X}_i$? Following statistical mechanics, we reach

$$P(\mathcal{X}_i | U_K) = \frac{Z_i}{Z} = \frac{\int_{x \in \mathcal{X}_i} \exp(-\frac{1}{RT} \hat{f}_i(x)) dx}{\sum_{j=1}^K \int_{x \in \mathcal{X}_j} \exp(-\frac{1}{RT} \hat{f}_j(x)) dx}$$

where Z_i and Z are local and global partition functions, respectively; and $\hat{f}_i(x)$ is the final Kriging regressor. Different regions are assumed to be mutually exclusive. The integrals are calculated by Monte Carlo sampling.

Another approach is to replace $\frac{1}{RT}$ above with ρ_i , the final ρ for temperature control in $\mathcal{X}_i$. In practice we did not find significant performance difference between the two approaches, partly due to the fact that final ρ_i in various clusters / regions reached similar values for the same protein complex.

2.6.3 $P(U_K)$

What is chance that the native structure is within the union of the initial regions, i.e., some initial region is near-native? The way to calculate $P(U)$ is very similar to that in uncertainty quantification. Specifically, 100,000 structures are sampled following the posterior of each region $\mathcal{X}_i$, evaluated for binding energy using the Kriging regressor $\hat{f}_i(x)$, and estimated with iRMSD using the binding affinity predictor formula inversely. We empirically calculate the portion q_i in which sample iRMSD values are above 4Å. Assuming the independence among regions, we reach $P(U_K) = 1 - \prod_{i=1}^K q_i$.

2.7 Data sets

We use a comprehensive protein docking benchmark set 4.0 Hwang *et al.* (2010) of 176 protein pairs that diversely and representatively cover sequence and structure space, interaction types, and docking difficulty. We split them into a training set, test sets **a** and **b** with stratified sampling to preserve the composition of difficulty levels in each set. The “training” set is not used for tuning BAL parameters (Sec. 2.2). Rather, it is just for training energy model ($y(x; w)$ in Sec. 2.4) and conformational-change extent prediction (τ_R in Sec. 2.3.3). The training and test **a** sets contain 50 and 26 pairs with known K_d values ((Kastritis and Bonvin, 2010)), respectively. And the test set **b** contain 100 pairs with K_d values predicted from sequence alone (Yugandhar and Gromiha, 2014).

We also use a smaller yet more challenging CAPRI set of 15 recent CAPRI targets (Chen *et al.*, 2017). Unlike the benchmark set for unbound docking, the CAPRI set contains 11 cases of homology docking, 8 of which start with just sequences for both proteins and demand homology models of structures before protein docking. Compared to the benchmark test set of 86 (68%), 22 (18%) and 18 (14%) cases classified rigid, medium, and flexible, respectively; the corresponding statistics for the CAPRI set are 4 (27%), 5 (33%) and 6 (40%), respectively. Their K_d values are also predicted from sequence alone.

The complete lists of the benchmark sets and the CAPRI set, with difficulty classification, are provided in Sec. 2.5 of the **SI**.

For each protein pair, we use 10 distinct encounter complexes as starting structures ($K = 10$). As reported previously (Chen *et al.*, 2017),

those for the benchmark sets are top-10 cluster representatives by ZDOCK, kindly provided by the Weng group; and those for the CAPRI set are top-10 models generated by the ZDOCK webserver.

3 Results and Discussion

3.1 Optimization and UQ Performance on Test Functions

We first tested our BAL algorithm on four non-convex test functions of various dimensions and compared it to particle swarm optimization (PSO) (Kennedy and Eberhart, 1995; Clerc and Kennedy, 2002), an advanced optimization algorithm behind a very successful protein-docking method SwarmDock (Moal and Bates, 2010). Detailed settings are provided in Sec. 2.7 of the **SI**.

For optimization we assess $\|\hat{x} - x^*\|$, the distance between the predicted and actual global optima, a measure of direct relevance to the quality of interest in protein docking – iRMSD. Compared to PSO, BAL made predictions that are, on average, closer to the global optima with smaller standard deviations; and the improvement margins increased with the increasing dimensions (Table S8 in the Supplementary Material).

For UQ we assess r_{90} , the distance upper bound of 90% confidence, i.e., $P(\|\hat{x} - x^*\| \leq r_{90}) = 1 - \sigma = 90\%$. The metrics to assess r_{90} include η , the relative error ($\eta = \frac{r_{90}}{\|\hat{x} - x^*\|} - 1$); and $\hat{P}$, the portion of the confidence intervals from 100 runs that actually encompass the corresponding global optimum. We found in Table S9 that our confidence intervals are usually tight judging from η and they contain the global optima with portions $\hat{P}$ close to 90%, the desired confidence level. The portions agreed less with the desired confidence level for some functions as the dimensionality increase, which suggests the challenge of optimization and UQ in higher dimensions.

3.2 Evaluation of Our Energy Models

We compare the performances of three machine learning models over the training and test sets for energy model (Sec. 2.4.2). As no actual binding affinities of encounter complexes are available, we estimated the iRMSD values based on the random forest model’s binding energy prediction (Sec. 2.4.1) and compared them to the actual iRMSD (of native interfaces) using RMSE for absolute error. Random forest gave the best performances thus used as the energy model $y(x)$ hereinafter. Specifically, performances are split to encounter complexes of varying quality (iRMSD) in Fig. 2A and Fig. 2B. The random forest model (blue bars) led to RMSE of 0.70Å (1.0Å) for the near-native samples in the training (test) set. The RMSEs increased slowly as iRMSD ≤ 10 Å and did sharply beyond (a region too far from the native for refinement).

We also assess how “funnel-like” the energy model is. We thus calculated for each protein pair the Spearman’s ranking coefficient ρ between the energy model and the actual iRMSD. The random forest of MM-GBSW features showed the highest ρ of 0.72 and 0.60 for the training and the test sets, respectively (Fig. 2C), albeit with large deviation across protein pairs.

We lastly assess the energy model’s ability to rank across protein pairs. Specifically, we estimated each native protein-complex’s binding energy by setting iRMSD to be zero in the energy model and compared the estimated and actual binding energy using RMSE and Pearson’s r in the supplemental Table S11. The random-forest energy model achieved 2.45 (4.78) Kcal/mol in RMSE and Pearson’s r of 0.79 (0.75) in binding energy $\Delta G = -RT \ln(K_d)$ for the training (test) set.

3.3 Docking Performance: Optimization

We show the improvements in PSO and BAL solution quality (measured by the decrease of iRMSD) against the starting ZDOCK solutions in Fig. S4 of the Supplementary Material. Speaking of the amount of improvement, BAL improved iRMSD by 1.2Å, 0.74Å, and 0.76Å for the training, test, and

Fig. 2: The root mean square error (RMSE) between the predicted and actual iRMSD for (A): Training set and (B): Test set. (C): The Spearman's ρ between predicted $y(\mathbf{x})$ and the real iRMSD for the training and test sets.

Dataset	$1 - \sigma$	0.99	0.95	0.90	0.85	0.80
Training	η	0.40 (0.23)	0.35 (0.18)	0.31 (0.17)	0.27 (0.15)	0.22 (0.10)
	$\hat{P}$	0.97	0.91	0.84	0.79	0.75
Test a	η	0.43 (0.26)	0.39 (0.21)	0.28 (0.16)	0.25 (0.13)	0.21 (0.09)
	$\hat{P}$	0.95	0.87	0.83	0.75	0.73
Test b	η	0.44 (0.22)	0.38 (0.16)	0.26 (0.14)	0.23 (0.10)	0.19 (0.08)
	$\hat{P}$	0.91	0.84	0.80	0.74	0.71
CAPRI	η	0.43 (0.20)	0.35 (0.13)	0.27 (0.11)	0.22 (0.10)	0.20 (0.09)
	$\hat{P}$	0.91	0.85	0.81	0.74	0.70

Table 1. Uncertainty quantification performances of BAL on protein docking based on η , the relative error in iRMSD; and $\hat{P}$, the portion of confidence intervals from 100 runs encompassing the global optima. For η , means (and standard deviations in parentheses) are reported.

CAPRI sets, respectively, outperforming PSO's corresponding measures of 0.82Å, 0.45Å, and 0.49Å. It also outperformed PSO for the more challenging near-native cases (note that BAL's iRMSD improvement for the near-native test set or CAPRI set was almost neutral). Speaking of the portion with improvement, BAL improved iRMSD in the near-native cases for 75%, 68%, and 73% of the training, test, and CAPRI sets, respectively; whereas the corresponding statistics for PSO were 59%, 50%, and 53%, respectively. More split statistics based on docking difficulty can be found in Fig. S5.

We also compared BAL and PSO solutions head-to-head over subsets of varying difficulty levels for protein docking (Fig.3). Overall, BAL's solutions are better (or significantly better by at least 0.5Å) than those of PSO for 70%-80% (31%-45%) of the cases, which was relatively insensitive to the docking difficulty level.

Both PSO and BAL use a single trajectory of 31 iterations and 630 samples for each region/cluster. Most time is on local structure minimization and energy evaluations using CHARMM (Brooks *et al.*, 2009). The BAL running time for optimization of each cluster thus almost linearly grows with the size of the protein pair (Fig. S8), ranging from 7 hours for a 200-residue complex to 13 hours for a 1700-residue one.

3.4 Docking Performance: Uncertainty Quantification

We next assess the UQ results for protein docking. Similar to that for test functions, we assess $r_{1-\sigma}$, the half length of $(1 - \sigma)$ confidence interval $[lb, ub]$, i.e., $P(lb \leq \text{iRMSD}(\hat{\mathbf{x}}, \mathbf{x}^*) \leq ub) = 1 - \sigma$. The metrics to assess $r_{1-\sigma}$ include η , the relative error ($\eta = |\frac{r_{1-\sigma}}{\text{iRMSD}(\hat{\mathbf{x}}, \mathbf{x}^*)} - 1|$); and $\hat{P}$, the portion of the confidence intervals that actually contain the corresponding native structure across all docking runs (10 for 10 models of each protein pair in each set).

Table 1 shows that the portions matched well with the confidence levels over all four data sets. Test set **b** and the CAPRI set did not have actual K_d values available and were thus impacted further by the uncertainty of K_d prediction, although the impact did not appear significant. There was a trade off between the confidence level and the length of the confidence interval, as narrower confidence intervals (with less η) corresponded to lower confidence levels. A balance seems to be at the 85% confidence level where the relative iRMSD uncertainty is around 25%.

3.5 Docking Performance: Quality Estimation

For quality estimation, two metrics are used for assessing the performance. The first is Spearman's ρ for ranking protein-docking predictions (structure models) for each pair. The second is the area under the Precision Recall Curve (AUPRC), for the binary classification of each prediction being near-native or not. Considering that the near-natives are minorities among all predictions, AUPRC is a more meaningful measure than the more common AUROC.

With these two metrics we assess four scoring functions on predictions $\hat{\mathbf{x}}_i$: (1) $\Delta E(\hat{\mathbf{x}}_i)$, the MM-GBSW binding energy, i.e., the sum of the 8 features; (2) our random-forest energy model $y(\hat{\mathbf{x}}_i)$; (3) $P(\mathcal{X}_i|U_K)$, the conditional probability that the i th prediction's region is near-native given that there is at least such one in the top K predictions; and (4) $P(\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4)$, the unconditional probability that the i th prediction is near-native.

For ranking assessment, from Fig. 4 we find that, whereas the original MM-GBSW model merely achieved merely 0.2 for Spearman's ρ , our energy model using the same 8 terms as features in random forest drastically improved the ranking performance with a Spearman's ρ around 0.6 for training, benchmark test, and CAPRI test sets. Furthermore, the confidence scores $P(\mathcal{X}_i|U_K)$ and $P(\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4)$ further improved ranking. In particular, the unconditional probability for a prediction to be near-native achieved around 0.70 in ρ even for the benchmark and the CAPRI test sets. Note that this probability, a confidence score on the prediction's near-nativeness, was derived from the posterior distribution of $\mathbf{x}^*$; thus it uses both enthalpic and entropic contributions.

For binary assessment on classifying the nativeness of predictions, the test set is split into **a**, 26 pairs with known K_d values and **b**, 100 with predicted ones. From Table 2 we conclude that the MM-GBSW energy model performed close to random (AUROC close to 0.5) and the random forest energy model using the same features drastically improved AUROC to around 0.8 and AUPRC 0.54 $\sim$ 0.62 across sets. Since AUROC is uninformative for highly imbalanced data (for instance, near-native predictions are 14% over all data sets), we focus on AUPRC. The next three probabilities from our BAL's confidence scores improved the AUPRC to nearly 0.80 for the training and above 0.60 for test sets. The additional uncertainty in K_d prediction from test sets **b** and CAPRI did not noticeably impact the performance compared to test set **a**.

3.6 Docking Results by BAL Optimization and Confidence-Score Ranking

We summarize our docking results (BAL predictions $\hat{\mathbf{x}}_i$ ranked by confidence scores on their nativeness $P(\text{iRMSD}(\hat{\mathbf{x}}_i, \mathbf{x}^*) \leq 4)$) in Table 3; and compare them to the ZDOCK starting results (ranked by cluster size roughly reflecting entropy) and the PSO refinement results (using the same energy model as BAL and ranked by the energy model). We use N_K to denote the number of targets with at least one near-native predictions in top K ; and F_K the fraction of such targets among all in a given set (training, benchmark test, or CAPRI test set). Compared to the ZDOCK

Fig. 3: The percentage of BAL predictions with iRMSD improvement against PSO for A. the training set, B. the benchmark test set, and C. the CAPRI set. The darker gray portions correspond to significant improvement (over 0.5 Å in iRMSD) compared to corresponding PSO predictions.

Fig. 4: Ranking performance shown in the bar plot (with error bar in black) of Spearman's ρ for (A) Training set, (B) Test set and (C) CAPRI set, respectively. a,b,c,d in each figure correspond to the MM-GBSW model, our random-forest energy model, and confidence scores $P(\mathcal{X}_i|U_K)$, and $P(\text{iRMSD}(\hat{x}_i, \mathbf{x}^*) \leq 4)$, respectively.

Dataset	Assessment	$\Delta E(\hat{x}_i)$	$y(\hat{x}_i)$	$P(\mathcal{X}_i U_K)$	$P(\text{iRMSD}(\hat{x}_i, \mathbf{x}^*) \leq 4 U_K)$	$P(\text{iRMSD}(\hat{x}_i, \mathbf{x}^*) \leq 4)$
Training	AUROC	0.489	0.806	0.903	0.944	0.967
	AUPRC	0.241	0.624	0.684	0.771	0.796
Testing a	AUROC	0.460	0.810	0.892	0.929	0.939
	AUPRC	0.199	0.550	0.592	0.613	0.634
Testing b	AUROC	0.490	0.789	0.847	0.898	0.927
	AUPRC	0.203	0.540	0.571	0.609	0.615
CAPRI	AUROC	0.491	0.771	0.844	0.893	0.919
	AUPRC	0.214	0.561	0.600	0.610	0.614

Table 2. Binary quality estimation for our 5 scoring functions on a prediction $\hat{x}_i$ being near-native or not: MM-GBSW energy model $\Delta E(\hat{x}_i)$, random-forest energy model $y(\hat{x}_i)$, and 3 BAL-determined probabilities that a region/cluster $\mathcal{X}_i$ is near-native given a native-containing list, the prediction $\hat{x}_i$ in that region is near-native given a native-containing list, or $\hat{x}_i$ is near-native.

starting results and PSO refinements, BAL has improved the portion of acceptable targets with top 3 predictions from 23% and 26%, respectively, to 32% for the benchmark test set. Similar improvements were found for the CAPRI set. The portion for top 10 from BAL reached 40% compared to ZDOCK's 33% over the benchmark test set. Notice that BAL only refined top 10 starting results from ZDOCK thus this improvement was purely from optimization (no ranking effect).

3.7 Energy Landscapes and Association Pathways

We lastly investigate energy landscapes during BAL sampling. Our kriging regressor $\hat{f}(\mathbf{x})$ is an unbiased estimator and works even better than the noisy observations from the random-forest energy model $y(\mathbf{x})$. Energy landscapes are visualized for 37 near-native regions for the benchmark test set (Fig. S6 in the Supplementary Material) and for the CAPRI set (Fig. S7) (non-rigid cases only). Two examples are shown in Fig. 5, depicting a (multiple) funnel-like energy landscape with a clear association paths from the starting to the end or the native complex along gradient descents. Three more supplemental videos are provided to visualize the BAL sampling trajectories using protein structures.

4 Conclusions

We present the first uncertainty quantification (UQ) study for protein docking. This is accomplished by a rigorous Bayesian framework that actively samples a noisy and expensive black-box function (i.e., collecting data D) while updating a posterior distribution $p(\mathbf{x}^*|D)$ directly over the

Fig. 5: The estimated energy landscapes along the first two principal components (PC) for two medium-difficulty docking cases with near-native starting models. Black dots are the samples. The grey triangle is the estimated end structure and the grey star is the true native structure. The starting structure is a thicker black dot at the origin. All the energy values are in the unit of RT and relative to the lowest sample energy value within each region.

unknown global optimum $\mathbf{x}^*$. The iterative feedback between Thompson sampling and posterior updating is linked by a Boltzmann distribution with adaptive annealing schedule and non-parametric Kriging regressor. The inverse uncertainty quantification on the location of the global optimum can easily forward-propagate for the uncertainty quantification of any quality of interest as a function of the global optimum, including the interface RMSD that measures dissimilarity between protein-docking solutions and native structures.

We demonstrate the superb performances of Bayesian active learning (BAL) on a protein docking benchmark set as well as a CAPRI set full of homology docking. Compared to the starting points from initial rigid docking as well as the refinement from PSO, BAL shows significant improvement, accomplishing a top-3 near-native prediction for about one-third of the benchmark and CAPRI sets. Its UQ results achieve tight uncertainty intervals whose radius is 25% of iRMSD with a 85% confidence level attested by empirical results. Moreover, its estimated probability of a prediction being near-native achieves an AUROC over 0.93 and AUPRC over 0.60 (more than 4 times over random classification).

	ZDOCK (Starting Point)			PSO			BAL		
Dataset (size)	N_3 (F_3)	N_5 (F_5)	N_{10} (F_{10})	N_3 (F_3)	N_5 (F_5)	N_{10} (F_{10})	N_3 (F_3)	N_5 (F_5)	N_{10} (F_{10})
Training (50)	11 (22%)	13 (26%)	17 (34%)	15 (30%)	17 (34%)	20 (40%)	19 (38%)	20 (40%)	22 (44%)
Test (126)	29 (23%)	33 (26%)	41 (33%)	33 (26%)	39 (31%)	45 (36%)	40 (32%)	42 (33%)	50 (40%)
CAPRI (15)	2 (13%)	3 (20%)	4 (27%)	3 (20%)	4 (27%)	5 (33%)	4 (27%)	5 (33%)	5 (33%)

Table 3. Summary of docking results measured by the number and the portion of targets in each set that have an acceptable near-native top 3, 5, or 10 prediction.

Besides the optimization and UQ algorithms, we for the first time represent the conformational space for protein docking as a flat Euclidean space spanned by complex normal modes blending flexible- and rigid-body motions and anticipating protein conformational changes, a homogeneous and isotropic space friendly to high-dimension optimization. We also construct a funnel-like energy model using machine learning to associate binding energies of encounter complexes sampled in docking with their iRMSD. These innovations also contribute to the excellent performances of BAL; and lead to direct visualization of binding energy funnels and protein association pathways in conformational degrees of freedom.

Acknowledgements

We thank Thom Vreven and Zhiping Weng for the ZDOCK rigid-docking decoys for the benchmark set, Haoran Chen for helping on cNMA, and Yuanfei Sun for proofreading the manuscript. Part of the CPU time was provided by the Texas A&M High Performance Research Computing.

Funding

This work was supported by the National Institutes of Health (R35GM124952) and the National Science Foundation (CCF-1546278).

Conflict of interest: None declared.

References

Agrawal, S. and Goyal, N. (2011). Analysis of Thompson Sampling for the multi-armed bandit problem. *arXiv:1111.1797 [cs]*.

Brooks, B. R. *et al.* (2009). CHARMM: the biomolecular simulation program. *Journal of Computational Chemistry*, **30**(10), 1545–1614.

Chapelle, O. and Li, L. (2011). An Empirical Evaluation of Thompson Sampling. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2249–2257. Curran Associates, Inc.

Chen, H. *et al.* (2017). Predicting protein conformational changes for unbound and homology docking: learning from intrinsic and induced flexibility. *Proteins: Structure, Function, and Bioinformatics*, **85**(3), 544–556.

Chilès, J.-P. and Delfiner, P. (2012). *Geostatistics: Modeling Spatial Uncertainty*, 2nd Edition.

Clerc, M. and Kennedy, J. (2002). The particle swarm-explosion, stability, and convergence in a multidimensional complex space. *IEEE transactions on Evolutionary Computation*, **6**(1), 58–73.

Der Kiureghian, A. and Ditlevsen, O. (2009). Aleatoric or epistemic? does it matter? *Structural Safety*, **31**(2), 105–112.

Gray, J. J. *et al.* (2003). Protein-protein docking with simultaneous optimization of rigid-body displacement and side-chain conformations. *Journal of molecular biology*, **331**(1), 281–299.

Györfi, L. *et al.* (2002). *A Distribution-Free Theory of Nonparametric Regression*. Springer Series in Statistics. Springer-Verlag, New York.

Henrández-Lobato, J. M. *et al.* (2014). Predictive Entropy Search for Efficient Global Optimization of Black-box Functions. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS'14*, pages 918–926, Cambridge, MA, USA. MIT Press.

Hwang, H. *et al.* (2010). Protein-Protein Docking Benchmark Version 4.0. *Proteins*, **78**(15), 3111–3114.

Ingber, L. (2000). Adaptive simulated annealing (ASA): lessons learned. *CoRR*, **cs.MS/0001018**.

Jiménez-García, B. *et al.* (2017). Lightdock: a new multi-scale approach to protein-protein docking. *Bioinformatics*, **34**(1), 49–55.

Kastritis, P. L. and Bonvin, A. M. J. J. (2010). Are Scoring Functions in Protein-Protein Docking Ready To Predict Interactomes? Clues from a Novel Binding Affinity Benchmark. *Journal of Proteome Research*, **9**(5), 2216–2225.

Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization, proceedings of IEEE international conference on neural networks (icnn'95) in.

Kozakov, D. *et al.* (2005). Optimal clustering for detecting near-native conformations in protein docking. *Biophysical journal*, **89**(2), 867–875.

Li, W. *et al.* (2016). Estimation of Uncertainties in the Global Distance Test (GDT_{ts}) for CASP Models. *PLOS ONE*, **11**(5), e0154786.

Marze, N. A. *et al.* (2018). Efficient Flexible Backbone Protein-Protein Docking for Challenging Targets. *Bioinformatics (Oxford, England)*.

Mirzaei, H. *et al.* (2015). Energy minimization on manifolds for docking flexible molecules. *Journal of chemical theory and computation*, **11**(3), 1063–1076.

Moal, I. H. and Bates, P. A. (2010). SwarmDock and the Use of Normal Modes in Protein-Protein Docking. *International Journal of Molecular Sciences*, **11**(10), 3623–3648.

Mosca, R. *et al.* (2013). Interactome3d: adding structural details to protein networks. *Nature methods*, **10**(1), 47.

Méndez, R. *et al.* (2005). Assessment of CAPRI predictions in rounds 3–5 shows progress in docking procedures. *Proteins: Structure, Function, and Bioinformatics*, **60**(2), 150–169.

Oliwa, T. and Shen, Y. (2015). cNMA: a framework of encounter complex-based normal mode analysis to model conformational changes in protein interactions. *Bioinformatics*, **31**(12), i151–i160.

Ortega, P. *et al.* (2012). A Nonparametric Conjugate Prior Distribution for the Maximizing Argument of a Noisy Function. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 3005–3013. Curran Associates, Inc.

Pfeifferberger, E. and Bates, P. A. (2018). Refinement of protein-protein complexes in contact map space with metadynamics simulations. *Proteins: Structure, Function, and Bioinformatics*.

Rasheed, M. *et al.* (2017). Statistical Framework for Uncertainty Quantification in Computational Molecular Modeling. *IEEE/ACM Trans Comput Biol Bioinform.*

Russo, D. *et al.* (2017). A Tutorial on Thompson Sampling. *arXiv:1707.02038 [cs]*.

Shahriari, B. *et al.* (2016). Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, **104**(1), 148–175.

Shen, Y. (2013). Improved flexible refinement of protein docking in capri rounds 22–27. *Proteins: Structure, Function, and Bioinformatics*, **81**(12), 2129–2136.

Shen, Y. *et al.* (2007). Optimizing noisy funnel-like functions on the euclidean group with applications to protein docking. In *Decision and Control, 2007 46th IEEE Conference on*, pages 4545–4550. IEEE.

Shen, Y. *et al.* (2008). Protein docking by the underestimation of free energy funnels in the space of encounter complexes. *PLoS computational biology*, **4**(10), e1000191.

Smith, G. R. and Sternberg, M. J. (2002). Prediction of protein-protein interactions by docking methods. *Current opinion in structural biology*, **12**(1), 28–35.

Vajda, S. and Kozakov, D. (2009). Convergence and combination of methods in protein-protein docking. *Current opinion in structural biology*, **19**(2), 164–170.

Villemonteix, J. *et al.* (2006). An informational approach to the global optimization of expensive-to-evaluate functions. *CoRR*, **abs/cs/0611143**.

Yugandhar, K. and Gromiha, M. M. (2014). Protein-protein binding affinity prediction from amino acid sequence. *Bioinformatics*, **30**(24), 3583–3589.