Published in final edited form as:

Conf Proc IEEE Eng Med Biol Soc. 2006 ; 1: 5330–5333.

Protein-protein docking with reduced potentials by exploiting multi-dimensional energy funnels

Ioannis Ch. Paschalidis[†], Yang Shen[‡], Pirooz Vakili[¶], and Sandor Vajda[§]

[†]Center for Information & Systems Eng., and Dept. of Manufacturing Eng., Boston University, 15 St. Mary's St., Brookline, MA 02446, e-mail: yannis@bu.edu, url: <http://ionia.bu.edu/>.

[‡]Center for Information & Systems Engineering, and Dept. of Manufacturing Eng., Boston University, e-mail: yangshen@bu.edu.

[¶]Center for Information & Systems Engineering, and Dept. of Manufacturing Eng., Boston University, e-mail: vakili@bu.edu.

[§]Department of Biomedical Engineering, Boston University, e-mail: vajda@bu.edu.

Abstract

We propose a new computational approach for protein docking exploiting energy funnels in the 6-dimensional space of translations and rotations of the ligand with respect to the receptor. Our approach consists of a series of translational and orientational moves of the ligand towards the receptor. Each move is performed using a global optimization method we have developed – the Semi-Definite Underestimation (SDU) method – which can exploit a funnel-like energy function. We compared our approach with Monte Carlo on a set of 10 protein complexes using two residue-level potentials. To achieve the same level of performance (produce a near-native $\leq 3\text{\AA}$ RMSD complex) our approach reduces energy evaluations by more than a factor of two, on average.

Keywords

Computational biology; Global optimization; Semi-definite programming; Molecular docking

I. INTRODUCTION

The genome-wide proteomics study fosters the need for understanding protein–protein interactions. Although X-ray crystallography has determined the structure of many protein complexes, the number of such complexes is low compared to the total number of structures in the Protein Data Bank (PDB). Consequently, determining the atomic coordinates of the complex computationally – a problem known as *protein docking* – is critical. In protein docking, one starts from the structures of the component proteins – the *receptor* and the *ligand* – which are assumed known.

Given the fact that the *native* complex adopts the lowest Gibbs energy, the problem can be cast as global optimization of the binding free energy. According to recent *CAPRI* protein-protein docking experiment results,¹ all successful methods are based on a multistage approach. Such an approach starts with a systematic coarse grain search, where Fast Fourier Transform (FFT) correlation techniques [1]) are widely used. Then promising conformational subsets are further explored by more intelligent optimization algorithms in higher resolution. Such “refinement”-stage algorithms include Monte Carlo minimization with simultaneous rigid body and side-

¹<http://capri.ebi.ac.uk/>.

chain optimization [2], and pseudo-Brownian rigid body docking followed by Biased Probability Monte Carlo minimization of the ligand interacting side chains [3].

As in protein folding, the protein binding energy landscape is believed to contain extremely rugged funnels [4]. This concept is more or less computationally validated in the one-dimensional dissimilarity space of RMSD (Root Mean Square Distance). Further, it has a physical explanation. Generally, the energy includes forces that act in different space-scales, resulting in multi-frequency behavior, and leading to a huge number of local minima caused by high-frequency terms (such as van der Waals). This suggests that even in a relatively small region, a local search algorithm can be of low efficiency and trapped easily in kinetic moves [5]. But at local minima where the solute-solute and solute-solvent interfaces are equally well packed without overlaps, the intermolecular van der Waals interactions in the bound state are largely balanced by solute-solvent interactions in the free state. Therefore, at such local minima the free energy surface is essentially determined by the “smooth” free energy component, which exhibits a funnel-like shape [4]. However, the funnel-like shape has not been sufficiently explored in the higher-dimensional conformational space. This is exactly what our work sets out to achieve.

In this paper we present a stochastic global optimization algorithm for protein docking. It explores the energy funnels in the space of rigid motion and leads a rigid motion pathway toward the global minimum by exploiting such funnels.

II. MATERIALS AND METHODS

A. Reduced Potentials

To better understand the structure of funnels in the orientational subspace, and remove complicating factors present in the full version of the problem, we will work with two models. These models are Gō-type models [6] and have been extensively used for understanding protein folding. They classify all possible contacts as “native” (i.e., present in the native structure) or “non-native”. A set of potentials are developed so that native contacts are favorable while non-native are not. The major advantage of such reduced models is that they yield a smoother energy landscape.

In the simplified system, both proteins (receptor and ligand) are represented at residue level. The center of a residue is chosen as the position of the C_β atom except for Glycine where we use the C_α atom since its side-chain consists of only a hydrogen atom. The receptor is held fixed and the ligand is freely moving as a rigid body. Then, the conformation of the ligand, or more specifically, the positions of the ligand’s residues become a function of conformational variables which we denote by the vector $\mathbf{x}$. Let $\mathcal{R} = \{\mathbf{r}_i; i = 1, \dots, I\}$ denote the set of positions for the receptor residues where $\mathbf{r}_i$ is the center of receptor residue i and $\mathcal{L} = \{\mathbf{r}_j(\mathbf{x}); j = 1, \dots, J\}$ denote the set of positions for the ligand residues where similarly $\mathbf{r}_j(\mathbf{x})$ is the center of ligand residue j . Let $\mathbf{x}^*$ correspond to the native complex and define as “native contacts” all pairs (i, j) where i corresponds to the i th receptor residue, j corresponds to the j th ligand residue, and $d_l \leq \|\mathbf{r}_i - \mathbf{r}_j(\mathbf{x}^*)\| \leq d_u$ for some appropriate distances d_l, d_u (e.g., $d_l = 3.8 \text{ \AA}$ and $d_u = 6.5 \text{ \AA}$). Namely, we call “native” all contacts where the corresponding residues are close enough in the native complex. Denote by $\mathcal{N}$ the set of all these native contacts (i, j) and set $d_{ij}(\mathbf{x}) = \|\mathbf{r}_i - \mathbf{r}_j(\mathbf{x})\|$, where $d_{ij}(\mathbf{x}^*)$ will denote distances in the native complex. Measuring all distances in $\text{\AA}$, we consider the following two potentials:

Reduced Potential 1— $F_1(\mathbf{x}) = \sum_{(i,j) \in \mathcal{N}} (d_{ij}(\mathbf{x}) - d_{ij}(\mathbf{x}^*))^2 - 100$. Essentially, there are a number of “springs” connecting the residues of the receptor and the ligand at the interface. These springs have different equilibrium length $d_{ij}(\mathbf{x}^*)$ but the stretch force constants are not differentiated. For physical feasibility we also use the penalty term of space exclusions: 10

$\sum_{i,j} \mathbf{1}\{d_{ij}(x) < d_l\}$, where $\mathbf{1}\{\cdot\}$ is an indicator function equal to 1 only if the corresponding condition is satisfied and 0 otherwise. This term is defined for every reduced potential and would not be repeated in the next potentials. Notice that the potential is expected to exhibit relatively smooth funnel-like shape but not trivially quadratic in x since $d_{ij}(x)$ is nonlinear in x .

Reduced Potential 2—We maintain the penalty for space clashes but replace the native interactions by the term $F_2(x) = 10 \sum_{(i,j) \in E_{ij}} \mathbf{1}\{d_l \leq d_{ij}(x) \leq d_u\} + \sum_{(i,j) \in E_{ij}} (d_{ij}(x) - d_u) \mathbf{1}\{d_{ij}(x) > d_u\}$. Namely, we consider a strong attraction force for native contacts in the range (d_l, d_u) and a linear force for longer range contacts. The native interactions are weighted with the Miyazawa-Jernigan residue-residue contact energies [7], denoted by E_{ij} for residue pair (i, j) .

The purpose of local minimization for these reduced potentials is to remove space clashes at the residue level. This is simply achieved by incrementally pulling the ligand away from the receptor along the line segment connecting the two centers of mass, until the penalty term becomes zero. The incremental stepsize we used was 0.2 Å.

B. The SDU Method

We have developed a stochastic optimization method called *Semi-Definite Underestimator (SDU)* algorithm for funnel-like shaped scoring functions. The method has similarities with the CGU method [5], applied in protein folding, which uses canonical quadratic underestimation to approximate the energy function. However, the rather restricted choice of the underestimator results in problematic performance for challenging protein docking problems. We follow a similar strategy and work on the envelope surface spanned by the set of local minima. This surface inherits the smooth behavior of the low-frequency energy terms. More specifically, we generate a moderate number of local minima and construct a *general* convex quadratic function that forms the tightest underestimator of all of them. This quadratic function suggests the location of the energy minimum. We use this information to iteratively refine our search.

Constructing an underestimator—Let us denote by $f: \mathbb{R}^n \rightarrow \mathbb{R}$ the free energy function we seek to minimize and assume we have obtained a set of K local minima $\phi^1, \dots, \phi^K$ of $f(\cdot)$. We are interested in constructing a convex quadratic function $U(\phi)$ satisfying $U(\phi^i) \leq f(\phi^i)$, for all $i = 1, \dots, K$, that is, a function that underestimates $f(\cdot)$ at all local minima ϕ^i , $i = 1, \dots, K$. More specifically, $U(\phi) \triangleq \phi'Q\phi + b'\phi + c$, where $Q \in \mathbb{R}^{n \times n}$, is a positive semi-definite matrix, $b \in \mathbb{R}^n$, and c is a scalar. The positive semi-definiteness of Q guarantees the convexity of $U(\cdot)$.

In [8] we showed that the problem of finding the tightest possible such underestimator $U(\cdot)$ can be formulated as a *Semi-Definite Programming (SDP)* problem. These are nonlinear problems that possess special structure and can be solved efficiently using interior-point methods (in polynomial time in the size of the input).

Biased sampling—Suppose we are seeking the global minimum of $f(\cdot)$ in some region $\mathcal{B}$ and have obtained an under-estimator $U(\cdot)$ as described above. Depending on the samples we used, and assuming that the constructed underestimator reflects the general structure of $f(\cdot)$, the minimum of $U(\cdot)$, say ϕ^P , is in the vicinity of the global minimum of $f(\cdot)$. We will be referring to ϕ^P as the *predictive conformation*. Our sampling method generates points so that the ones close to ϕ^P are more likely to be selected while points with high enough energies are assigned small probabilities (see [8]).

The SDU algorithm (outline) [8]—SDU seeks a global minimum of $f(\cdot)$ in a region $\mathcal{B}$ and maintains a set $\mathcal{L}$ of interesting local minima obtained so far as well as the best such local minimum denoted by ϕ^G . It goes through a series of iterations, each consisting of an *exploration* step and an *underestimation* step. The exploration step generates points in the current search region using the sampling approach described above. Using each such sample as a starting point we perform local minimization of $f(\cdot)$ and update $\mathcal{L}$ to include these new local minima while discarding unfavorable local minima. Solving the SDP problem described earlier we derive a new underestimator that underestimates $f(\cdot)$ at all points in $\mathcal{L}$, which is then used to drive exploration in the next iteration. In [8] we have established that SDU probabilistically converges (i.e., with probability converging to 1) to the global minimum of $f(\cdot)$ when applied to a general class of funnel-shaped functions as the number of local minima used for underestimating $f(\cdot)$ grows.

C. Search Space and Our Strategy

For protein docking, we optimize the free energy over the 6-dimensional (6D) space, denoted by $SE(3)$, of translations and rotations of the ligand with respect to the receptor. $SE(3)$ is the semidirect product of two distinctive subspaces: $\mathbb{R}^3$ (translations) and $SO(3)$ (the rotation group). $\mathbb{R}^3$ is a Euclidean space without any curvatures. As to $SO(3)$, Euler first showed that $SO(3)$ is a 3-dimensional manifold and at least 5 parameters are required to represent it in an 1 – 1 global manner. One can also view optimization in $SO(3)$ as taking place in the vector space of 3×3 rotation matrices $\mathbf{R}$ where $\mathbf{R}^T\mathbf{R} = \mathbf{I}$ and $\det(\mathbf{R}) = 1$, where $\mathbf{I}$ is the identity matrix and $\det(\cdot)$ denotes the determinant. Various other parametrizations of $SO(3)$ are available, e.g., Euler angles, exponential coordinates, quaternions (see [9]).

The translational space $\mathbb{R}^3$ is Euclidean so applying SDU is more straightforward. We tested on unbound ligands for all complexes in the benchmark set of [10]. For each complex, we considered the bound orientation and pulled the ligand 7Å away along the line segment connecting the two centers of mass. We performed translational optimization while keeping the ligand's orientation fixed to the native one. Out of the 48 protein complexes in the benchmark set, SDU finds 19 predictions with $< 1\text{\AA}$ RMSD and 40 with $< 3\text{\AA}$ RMSD.

$SO(3)$, on the other hand, is a nonlinear manifold. It can be locally approximated by a Euclidean space, but, in such small regions the energy landscape is rugged and relatively flat to provide some guidance. As indicated in Sec. I, the protein-docking community has explored 1-dimensional energy funnels (in RMSD plots); yet, high-dimensional funnels in $SE(3)$ or $SO(3)$ have neither been validated nor explored. To overcome these difficulties, the strategy we adopt is to pursue optimization separately in each of the two subspaces – the translational subspace $\mathbb{R}^3$ and the orientational subspace $SO(3)$ – and then weave these moves into a coordinated movement of the receptor towards the ligand.

D. Parametrizations of the Rotation Group

To optimize these reduced potentials over $\mathbf{x} \in SE(3)$ the parametrization of $SO(3)$ is critical. The shortest curves – or minimal geodesics – on the $SO(3)$ manifold are of the form $\mathbf{R}(t) = \mathbf{R}_0 e^{\mathbf{\Omega}t}$, where $\mathbf{\Omega}$ is a 3×3 skew-symmetric matrix formed by the elements of some vector $\mathbf{\omega} \in \mathbb{R}^3$. In this curve, $\mathbf{R}_0 \in SO(3)$ is the initial point and $\mathbf{R}(1)$ is the final point. Hence, a parametrization of $SO(3)$ can be obtained by mapping rotation matrices $\mathbf{R} \in SO(3)$ to vectors $\mathbf{\omega} \in \mathbb{R}^3$ through the relationship $\mathbf{R} = \mathbf{R}_0 e^{\mathbf{\Omega}}$, where $\mathbf{\Omega}$ is a skew-symmetric matrix formed by the elements of $\mathbf{\omega}$. This parametrization produces deep and wide enough funnels of protein binding energies.

E. Coordinating Translational and Rotational Moves

The ultimate goal in refinement-stage docking is to start from some initial position and orientation of the ligand and move in the conformational space defined by the complexes in some conformational regions to form the minimum energy complex in the region. If the region we are working with contains the native complex, then our objective is to produce a high-quality ($< 1\text{--}3\text{\AA}$ RMSD) approximation.

Let us represent the energy function to be minimized as $f(\mathbf{r}, \boldsymbol{\omega})$ where the receptor is held fixed, $\mathbf{r} \in \mathbb{R}^3$ denotes the position of the ligand, and $\boldsymbol{\omega} \in \mathbb{R}^3$ maps to an orientation of the ligand in $SO(3)$ using the exponential parametrization described above. We use the following series of translational/orientational adjustments:

1. Orientational Adjustment—Start with a given starting structure characterized by $(\mathbf{r}^0, \boldsymbol{\omega}^0)$ and optimize the energy over $\boldsymbol{\omega}$ within some hypercube centered at $\boldsymbol{\omega}^0$ while keeping the translation vector fixed at $\mathbf{r}^0$. This optimization can be done using our SDU algorithm as $\boldsymbol{\omega}$ is unconstrained in a Euclidean space. Suppose this yields a new structure characterized by $(\mathbf{r}^0, \boldsymbol{\omega}^1)$; this is a structure that has oriented itself to minimize the energy potential.

2. Translational Adjustment—Start with $(\mathbf{r}^0, \boldsymbol{\omega}^1)$ and now optimize over $\mathbf{r}$ in some region centered at $\mathbf{r}^0$. That is, using the $\boldsymbol{\omega}^1$ orientation make a “step” towards the receptor in order to minimize the energy potential. Again, this can be done using our SDU algorithm.

3. Convergence Criterion—Repeat steps 1 and 2 above until no significant movement of the ligand is observed.

We note that even though each translational and orientational adjustment is in some small region of $\mathbb{R}^3$ and $SO(3)$, respectively, it is achieved by global optimization. That is, these moves are not local (in the sense of local optimization) and involve overcoming significant energy barriers.

III. RESULTS

We tested the proposed approach against a standard Monte Carlo minimization using the Metropolis acceptance criterion. The test set consists of 10 arbitrarily chosen bound protein complexes from the benchmark set of [10]: 5 enzyme-inhibitor complexes (1AVW, 1BRC, 1CSE, 2KAI and 2PTC) and 5 antigen-antibody or other ones (1A0O, 1AHW, 1AVZ, 1MLC and 2JEL). We place the ligand about 10 Å away from the receptor and orient it by rotating the native orientation by 0.8 radian on average (i.e., 46°) around an arbitrary axis. This results on average in ligands 15 – 20 Å away (binding site of the ligand, C_α) from the native one. Our initial search region consists of a 12 Å cube in the translational subspace and a 0.3 radian cube in the orientational subspace of local exponential parameters. The native structure x^* is beyond this region. Samples are uniformly generated in both of these regions until underestimators are constructed. Then biased sampling guided by the underestimators is adopted. Although function evaluation is negligible in these cases with reduced potentials, we should keep in mind that in docking actual proteins function evaluations (e.g., using CHARMM [11] or other complex potentials) dominate all other tasks. Hence, both algorithms take at most 1000 energy evaluations for each sample path and we collected 100 sample paths for each algorithm per protein complex.

We report the success rates η of both algorithms for each protein where “success” is defined based either on energy ($F(x) \leq 0.95 \times F(x^*)$, cf. Table I) or on similarity ($RMSD \leq 3 \text{ \AA}$, cf. Table II). K is the required number of sample paths (or, equivalently, energy evaluations in the

thousands) to guarantee a success with a probability of 95% ($K = \max\{\log_{(1-\eta)}(1-P), 1\}$). The last row of each table reports average results over all 10 complexes.

Based on the results above, we clearly see that SDU outperforms Monte Carlo. Using the more challenging reduced potential 2 and defining success based on energy, SDU needs on average 7,930 energy evaluations compared to 11,970 for Monte Carlo in order to reach the same level of performance; an efficiency gain of 33.8% (cf. Table I). If instead we define success based on similarity, the corresponding numbers are 5,900 for Monte Carlo and 2,830 for SDU, i.e., *SDU cuts energy evaluations by a more than a factor of two* (cf. Table II). It can also be seen that for antigen-antibody or other complexes the gap is even wider. In docking with more complex potentials, these comparisons reflect the relative computational requirements of the two algorithms since energy evaluations would dominate all other tasks.

This significant gain in efficiency is due to SDU's ability to exploit energy funnels, thus, zooming in more quickly in interesting regions of the energy landscape. We acknowledge that standard Monte Carlo minimization can be improved by simulated annealing-type modifications. Yet, we do not expect this to drastically change the qualitative advantage of SDU which benefits from the global structure of the energy landscape rather than relying on mostly local moves (as Monte Carlo does).

IV. CONCLUSION

We developed a new computational approach for protein docking exploiting energy funnels in the 6-dimensional space of translations and rotations of the ligand with respect to the receptor. The approach consists of a series of translational and orientational moves of the ligand towards the receptor. Each move is performed using the Semi-Definite Underestimation (SDU) method, introduced in [8], which can exploit funnel-like energy functions.

Using two residue-level potentials, we compared our approach with Monte Carlo on a set of 10 protein complexes, including enzyme-inhibitor, antigen-antibody, and other. To achieve the same level of performance (produce a near-native $\leq 3\text{\AA}$ RMSD complex) our approach reduces energy evaluations by more than a factor of two, on average. If instead the requirement is to produce a conformation with energy within 5% of the native, then the efficiency gain is 33.8%, on average.

Acknowledgements

Research partially supported by the NSF under grants DMI-0330171, ECS-0426453, CNS-0435312, DMI-0300359, EEC 0088073, by the NIH under grant R01-GM61867, and by the ARO under the ODDR&E MURI2001 Grant DAAD19-01-1-0465.

REFERENCES

1. Katchalski-Katzir E, Shariv I, Eisenstein M, Friesem AA, Aflalo C, Vakser I. Molecular surface recognition: Determination of geometric fit between proteins and their ligands by correlation techniques. *Proc. Natl. Acad. Sci. USA* 1992;vol. 89:2195–2199. [PubMed: 1549581]
2. Gray JJ, Moughon S, Wang C, Schueler-Furman O, Kuhlman B, Rohl CA, Baker D. Protein-protein docking with simultaneous optimization of rigid-body displacement and side-chain conformations. *J. Molecular Biology* 2003;vol. 331:281–299.
3. Fernandez-Recio J, Totrov M, Abagyan R. Identification of protein-protein interaction sites from docking energy landscapes. *J. Mol. Biol* 2004;vol. 335(no 3):843–865. [PubMed: 14687579]
4. Camacho C, Vajda S. Protein docking along smooth association pathways. *Proc. Natl. Acad. Sci. USA* 2001;vol. 98:10636–10641. [PubMed: 11517309]

5. Phillips, A.; Rosen, J.; Dill, K. ch. Convex Global Underestimation for Molecular Structure Prediction. In: Pardalos, PM., et al., editors. From Local to Global Optimization. Kluwer Academic Publishers; 2001. p. 1-18.
6. Taketomi H, Ueda Y, Gō N. Studies on protein folding, unfolding and fluctuations by computer simulation. I. the effect of specific amino acid sequence represented by specific inter-unit interactions. Int. J. Pept. Protein Res 1975;vol. 7:445–459. [PubMed: 1201909]
7. Miyazawa S, Jernigan RL. Residue-residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading. J. Mol. Biol 1996;vol. 256(no 3):623–644. [PubMed: 8604144]
8. Paschalidis, IC.; Shen, Y.; Vajda, S.; Vakili, P. A semi-definite programming-based underestimation method for global optimization in molecular docking. Proceedings of the 44th IEEE Conference on Decision and Control; Seville, Spain. 2005. p. 3675-3680.
9. Stuelpnagel J. On the parametrization of the three-dimensional rotation group. SIAM Review 1964;vol. 6(no 4):422–430.
10. Chen R, Mintseris J, Janin J, Z W. A protein-protein docking benchmark. Proteins 2003;vol. 52:88–91. [PubMed: 12784372]
11. Brooks B, Bruccoleri R, Olafson B, States D, Swaminathan S, M K. CHARMM: A program for macromolecular energy, minimization, and dynamics calculations. J Comput Chem 1983;vol. 4:187–217.

Fig. 1.

SDU forms the convex quadratic underestimator $U(x)$ of a funnel-like function $f(x)$ and focuses its exploration in the vicinity of $U(x)$'s minimum (using a sampling pdf proportional to $-U(x)$). Underestimation and focused exploration is repeated iteratively until convergence.

TABLE 1
RESULTS WITH REDUCED POTENTIALS 1 AND 2 BASED ON ENERGY.

Potential Algorithm Measure	Reduced 1			Reduced 2		
	MC η (%)	K	SDU η (%)	MC η (%)	K	SDU η (%)
1AVW	71	2.4	93	13	21.5	34
1BRC	89	1.4	98	19	14.2	25
1CSE	66	2.8	87	32	7.8	40
2KAI	82	1.7	93	24	10.9	23
2PTC	79	1.7	97	21	12.7	23
			Enzyme-Inhibitor Complexes			
			Antigen-Antibody or Other Complexes			
1A00	92	1.2	100	44	5.2	68
1AHW	58	3.5	92	19	14.2	20
1AVZ	61	3.1	83	34	7.2	62
1MLC	45	5.0	94	17	16.1	48
2JEL	88	1.4	96	26	9.9	28
Average	73.1	2.42	93.3	24.9	11.97	37.1
						7.93

NIH-PA Author Manuscript

NIH-PA Author Manuscript