Published in final edited form as:

Proc IEEE Conf Decis Control. 2012 December ; : 2983–2988. doi:10.1109/CDC.2012.6426267.

A New Approach to Rigid Body Minimization with Application to Molecular Docking*

Hanieh Mirzaei[†], Dima Kozakov[‡], Dmitri Beglov[‡], Ioannis Ch. Paschalidis[§], Sandor Vajda[‡], and Pirooz Vakili[¶]

[†]Division of Systems Eng., Boston University, hanieh@bu.edu

[‡]D. Kozakov, D. Beglov, and S. Vajda are with the Dept. of Biomedical Eng., Boston University, {midas, dbeglov, vajda}@bu.edu

[§]Dept. of Electrical & Computer Eng., and Division of Systems Eng., Boston University, yannisp@bu.edu

Abstract

Our work is motivated by energy minimization in the space of rigid affine transformations of macromolecules, an essential step in computational protein-protein docking. We introduce a novel representation of rigid body motion that leads to a natural formulation of the energy minimization problem as an optimization on the $SO(3) \times \mathbb{R}^3$ manifold, rather than the commonly used $SE(3)$. The new representation avoids the complications associated with optimization on the $SE(3)$ manifold and provides additional flexibilities for optimization not available in that formulation. The approach is applicable to general rigid body minimization problems. Our computational results for a local optimization algorithm developed based on the new approach show that it is about an order of magnitude faster than a state of art local minimization algorithms for computational protein-protein docking.

I. INTRODUCTION

Consider the minimization of a cost function where the cost depends on the location of a rigid body in the 3 dimensional space, $\mathbb{R}^3$. This is a prototype of a rigid body optimization problem in a number of applications such as workpiece or camera localization or calibration (see, e.g., [1], [2], [3]) and, as we will describe in this paper, molecular docking. More general versions may include more than one rigid body, and/or consider bodies that have rigid parts hinged together. In the single rigid body case, the optimization is naturally defined on the space of translations and rotations of the rigid body. The space of rigid body transformations in $\mathbb{R}^3$ is commonly represented by the Special Euclidean group (or simply the Euclidean group) $SE(3)$, a Lie group (see, e.g. [4]). With this representation, the rigid body optimization is formulated as a manifold optimization problem. The advantage of this

*Research supported in part by NIH grants 1-R01-GM093147-01 and GM061867.

[¶] Corresponding author. Dept. of Mechanical Eng. and Division of Systems Eng., Boston University, vakili@bu.edu.

formulation, compared to available alternatives, is that the optimization is defined on a low-dimensional space.

Many standard optimization algorithms on Euclidean spaces generalize to (Riemannian) manifolds (see, e.g., [5]). Consider, for example, the steepest descent algorithm. Once the steepest descent direction is identified, the generalization of line search for Euclidean spaces corresponds to searching along a geodesic on the manifold. As has been pointed out (see, e.g., [5], [6]), the efficiency of such generalizations partially depends on the ease with which we can compute geodesics of the manifold. For some manifolds that have a group structure compatible with their Riemannian structure, the geodesics are one-parameter subgroups and can be easily computed. This is, for example, the case for the Special Orthogonal group $SO(3)$.

$SE(3)$ is a semi-product of the Lie groups of rotation, $SO(3)$, and translations, $\mathbb{R}^3$. For $SE(3)$ it turns out that there is a “mismatch” between the natural Riemannian metric on $SO(3) \times \mathbb{R}^3$, as a direct product of $SO(3)$ and $\mathbb{R}^3$ manifolds, and the group structure of $SE(3)$. This is often highlighted by stating that $SE(3)$ does not have a natural bi-invariant Riemannian metric. Another implication of this mismatch is that geodesics of $SE(3)$ are no longer one-parameter subgroups of $SE(3)$. These make direct generalizations of standard Euclidean optimization algorithms to $SE(3)$ non-trivial (for an informative discussion, see [2]). Various remedies have been suggested to address this difficulty. For example, [2] proposes a cyclic coordinate descent approach where the optimization algorithm cycles through steps of optimizations over the $SO(3)$ component and the $\mathbb{R}^3$ component iteratively, taking full advantage of the group structures of the component manifolds. By contrast, [3] discards the group structure of $SE(3)$ altogether and considers optimization over $SO(3) \times \mathbb{R}^3$ as a Riemannian manifold only.

The first contribution of this paper is to present an alternative group of rigid body transformations that corresponds to the Lie group $SO(3) \times \mathbb{R}^3$, i.e., the direct product of the Lie groups $SO(3)$ and $\mathbb{R}^3$. Given this formulation, the rigid body optimization problem can be naturally defined as an optimization on $SO(3) \times \mathbb{R}^3$. As a result, the difficulties associated with optimization on $SE(3)$ can be fully avoided. The new formulation provides an additional flexibility to the user, namely choosing the initial center of rotation. This flexibility, as we point out in the paper, can be used to improve the performance of the optimization algorithm. It is worth noting that our new formulation is not tied to any particular application and is generally applicable to all rigid body optimization problems.

The second contribution of the paper is to use our formulation and a manifold optimization approach to address local optimization problems in molecular docking. A brief introduction to molecular docking and local optimization problems in that context are given in the paper. Here, we point out that a manifold optimization approach has not been used before for molecular docking. Our goal in this part was to develop optimization algorithms that are more efficient than the state of the art algorithms currently in use. Our choices in manifold optimization, as will be described in the paper, were driven by this practical consideration.

Our experimental results show that our algorithm substantially outperforms alternatives for protein-protein docking problems.

The rest of the paper is organized as follows. We introduce our alternative group of rigid body transformations in Section II. Section III gives a short introduction to molecular docking and provides the broader context for the local optimization we consider in this paper. In Section IV we define the local optimization problem, describe a set of choices that we have made, and present our local optimization algorithm. Section V includes our experimental results for protein-protein docking. We conclude in section VI.

II. AN ALTERNATIVE GROUP OF RIGID BODY TRANSFORMATIONS

As stated before, the common approach to representing rigid body transformations/motion is via the Euclidean group $SE(3)$, see, e.g., [7], [8]. In this section, we introduce an alternative representation of rigid body transformations. Before introducing the alternative representation, we give a brief review of the Euclidean group $SE(3)$.

A. Euclidean group $SE(3)$

Let

$$SO(3) = \{R \in \mathbb{R}^{3 \times 3}; R^T R = I; \det(R) = 1\}$$

denote the group of orientation-preserving rotations on $\mathbb{R}^3$ and consider the set $SO(3) \times \mathbb{R}^3 = \{(R, t); R \in SO(3), t \in \mathbb{R}^3\}$. Let $g = (R, t)$ and $g' = (R', t')$ be two elements of $SO(3) \times \mathbb{R}^3$. Then, the group “multiplication” of $SE(3)$ is defined by

$$g' g = (R' R, R' t + t').$$

The action of $SE(3)$ on $\mathbb{R}^3$ is defined as follows (see, e.g., [4], section 2.4). For $g \in SE(3)$, let $g : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be defined by

$$g(q) = Rq + t.$$

This mapping defines an action of the $SE(3)$ group on $\mathbb{R}^3$ since

$$g' \circ g(q) = g'(g(q)) = R'(Rq + t) + t' = R'Rq + R't + t' = g'g(q).$$

Therefore

$$g' \circ g = g'g$$

where $\circ$ denotes composition of functions.

Remark—We note that there is not a unique way to associate (R, t) with a rigid body motion. The unspecified element is the *center of rotation*. In the case of the Euclidean group $SE(3)$, the center of rotation is the origin of a fixed coordinate system if we adopt what is referred to as an active view (see [4], section 2.5).

It can be easily verified that this action corresponds to rigid body transformations of $\mathfrak{R}^3$ (in the stricter sense defined in [7], ch 2). Moreover, $SE(3)$ has a group representation known as the homogeneous representation where

$$(R, t) \rightarrow \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}.$$

The group multiplication of $SE(3)$ corresponds to matrix multiplication of 4×4 matrices with the above structure.

We now introduce the alternative representation.

B. Direct product group $SO(3) \times \mathfrak{R}^3$

The direct product “multiplication” on $SO(3) \times \mathfrak{R}^3$ is naturally defined by

$$g' * g = (R' R, t' + t).$$

We use $*$ to denote the direct product multiplication and to distinguish it from the multiplication of $SE(3)$.

The novel element of our representation is the action we associate with this group. We define this action on $\mathfrak{R}^3 \times \mathfrak{R}^3$ as follows. For $g \in SO(3) \times \mathfrak{R}^3$, let $g : \mathfrak{R}^3 \times \mathfrak{R}^3 \rightarrow \mathfrak{R}^3 \times \mathfrak{R}^3$ be defined by

$$g(q, p) = (R(q - p) + p + t, p + t), \\ (q, p \in \mathfrak{R}^3).$$

In words, the action of g on the first component $q \in \mathfrak{R}^3$ is to rotate q according to the rotation matrix R but with the center of rotation p and translate it by t . The action of g on the second component simply translates the point p by t . Equivalently, we can think that the action on the second component is of the same type as the action on the first component since $R(p - p) + p + t = p + t$. The following is an immediate result.

Proposition 1—The above transformation defines an action of the group $SO(3) \times \mathfrak{R}^3$ on $\mathfrak{R}^3 \times \mathfrak{R}^3$.

Proof: Let $g = (R, t)$ and $g' = (R', t')$. Then,

$$\begin{aligned}
g'(g(q, p)) &= g'(R(q - p) + p + t, p + t) \\
&= (R'(R(q - p) + p + t - (p + t)) \\
&\quad + p + t + t', p + t + t') \\
&= (R'R(q - p) + p + t + t', p + t + t') \\
&= (g' * g)(q, p).
\end{aligned}$$

We define the following representation for the direct product group $SO(3) \times \mathbb{R}^3$. Let $(q, p, 1)^T \in \mathbb{R}^7$. Then, let

$$(R, t) \rightarrow \begin{bmatrix} R & I - R & t \\ 0 & I & t \\ 0 & 0 & 1 \end{bmatrix}.$$

The group multiplication of $SE(3)$ corresponds to matrix multiplication of 7×7 matrices with the above structure.

We next show that, for any center of rotation $p \in \mathbb{R}^3$, the action of $SO(3) \times \mathbb{R}^3$ on $\mathbb{R}^3 \times \mathbb{R}^3$ is a rigid body transformation of the first component $\mathbb{R}^3$.

Let $\pi_i: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ ($i=1, 2$) be projections on the first and second coordinate ($\pi_1(q, p) = q$, $\pi_2(q, p) = p$). For any fixed $p \in \mathbb{R}^3$, let

$$g_p: \mathbb{R}^3 \rightarrow \mathbb{R}^3 \times \mathbb{R}^3$$

be defined by $g_p(q) = g(q, p)$. Then, we have

Proposition 2—For any p

$$\pi_1 \circ g_p: \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

is a rigid body transformation of $\mathbb{R}^3$.

Proof: Fix $p \in \mathbb{R}^3$. Let $q, q' \in \mathbb{R}^3$, then

$$\begin{aligned}
\|\pi_1 \circ g_p(q) - \pi_1 \circ g_p(q')\| &= \\
\|R(q - p) + p + t - (R(q' - p) + p + t)\| &= \\
\|R(q - q')\| &= \|q - q'\|.
\end{aligned}$$

The last equality is due to the fact that R is a rotation matrix. Following the definition in [7], we also need to show that $\pi_1 \circ g_p$ is orientation-preserving. In other words, it sends right-handed coordinate frames to right-handed coordinate frames. We show that the action of π_1

$\circ g_p$ on vectors in $\mathbb{R}^3$ is the same as the action of $SE(3)$ on such vectors. Therefore, the result follows from the fact that $SE(3)$ is orientation-preserving ([7], Proposition 2.7).

Let $q, q' \in \mathbb{R}^3$ as above. Then, under the $\pi_1 \circ g_p$ transformation, the vector $q' - q$ is transformed into the vector $\pi_1 \circ g_p(q') - \pi_1 \circ g_p(q)$. We showed above that the latter is equal to $R(q - q')$. Under $SE(3)$ transformation and $g = (R, t)$ $q' - q$ is transformed into the vector $Rq' + t - (Rq + t) = R(q' - q)$. Hence the proof that $\pi_1 \circ g_p$ is orientation-preserving.

A critical feature of this formulation of rigid body transformations is that, by contrast to the $SE(3)$ formulation, translational moves and rotational moves are decoupled. For example, let $g = (I, t)$, i.e., translation only by t and $g' = (R, 0)$, i.e., rotation only by R . Then it can be easily seen that in $SE(3)$

$$g g' \neq g' g$$

where as in $SO(3) \times \mathbb{R}^3$

$$g * g' = g' * g.$$

The above can be verified by considering the homogeneous representation of g and g' in $SE(3)$ and $SO(3) \times \mathbb{R}^3$ respectively.

The rigid body motion that we associate with $SO(3) \times \mathbb{R}^3$ is a natural motion in the context of the molecular docking problem that we consider next. Furthermore, since the group $SO(3) \times \mathbb{R}^3$ is a direct product of $SO(3)$ and $\mathbb{R}^3$ both as groups and as Riemannian manifolds, there is no mismatch between the group and the natural Riemannian structures and we do not face the complications that are associated with $SE(3)$ rigid body transformations.

Furthermore, as we mentioned earlier, in the $SO(3) \times \mathbb{R}^3$ formulation the user can choose the initial center of rotation. This gives a valuable flexibility to the user to better match the moves of the optimization with the dynamics of molecular interactions.

III. MOLECULAR DOCKING

In this section, we briefly review molecular docking and identify the local optimization problem that we consider later in the paper.

Predictive molecular docking is a fundamental and challenging problem in computational structural biology. Given two component molecules, termed *receptor* and *ligand*, the goal is to determine the most likely structure of a receptor-ligand complex by minimizing an energy-like target function. Elucidating the roles and structures of protein-protein complexes is crucial for a better understanding of processes such as metabolic control, signal transduction, immune response, and gene regulation.

The challenge for predictive docking is to start with the coordinates of the unbound component molecules and to obtain computationally a model of the bound complex [9], [10], [11]. One of the component molecules, usually the larger, will be considered as the receptor, and the other the ligand.

All successful state-of-the-art protein docking methods employ a so called multistage approach. They begin with large sampling of the conformational space with somewhat crude energy functions, then filter the resulting samples, and finally go through a refinement stage starting from filtered conformations. Independently of the algorithm used for sampling the conformational space, virtually all docking algorithms also include some type of local continuous minimization of the energy function in order to remove steric clashes and to obtain more reliable energy values [11]. For example, by fixing the position and orientation of the receptor and moving the ligand as a rigid body, the refinement stage can be viewed as an optimization of the energy function $E(x)$ where x denotes a rigid body movement of the ligand. E has a large number of local minima and the energy landscape is extremely rugged. Therefore, local information such as the gradient of E with respect to x is not a reliable guide for global optimization. A more global assessment of the energy landscape is needed for this purpose. We have developed a stochastic search algorithm called Semi-Definite Underestimation (SDU) that uses quadratic underestimators of the energy function to guide the search (see, [12], [13], [14]). However, due to the extreme ruggedness of the energy landscape, the energy of a sample x , $E(x)$, is not very informative about the landscape in the vicinity of x . The role of local minimization is to locate the minimum energy configuration x^* in the vicinity of x . The minimum energy, $E(x^*)$, is then used in an exterior global optimization loop such as SDU. Hence, the refinement stage involves repeated use of local minimization. This local minimization problem is what we address in the next section. Local information such as gradient of the energy function is informative at this stage and will be used in the algorithm we describe in the next section.

IV. LOCAL OPTIMIZATION ALGORITHM

We begin with defining more formally the energy function and the optimization problem.

A. Energy function

The receptor and ligand are specified by the coordinates of their individual atoms with respect to a fixed coordinate frame. Let m_r and m_l denote, respectively, the number of atoms of the receptor and ligand. Then, $\mathbf{q}^r = (q_1, \dots, q_{m_r}) \in \mathbb{R}^{3 \times m_r}$ and $\mathbf{q}^l = (q_1, \dots, q_{m_l}) \in \mathbb{R}^{3 \times m_l}$ correspond to specific configurations of the receptor and ligand.

The energy function, denoted by $E(\mathbf{q}^r, \mathbf{q}^l)$, reflects electrostatics, and van der Waals interactions as well as bonded energy terms. As stated earlier, we hold the receptor fixed and consider rigid movements of the ligand. Therefore, the energy function can be viewed as a function of the coordinates of the ligand only, i.e., $E(\mathbf{q}^l)$. Therefore, more abstractly, the energy function can be viewed as

$$E: \mathbb{R}^{3 \times m_l} \rightarrow \mathbb{R}.$$

For simplicity, we drop reference to the receptor and simply write $E(\mathbf{q})$ for $\mathbf{q} = (q_1, \dots, q_{m_l}) \in \mathbb{R}^{3 \times m_l}$. The energy function is a sum of a large number interaction terms that are explicitly given as functions of the coordinate of the atoms of the receptor and ligand. For details, see [13]. Here, we simply point out that the evaluation of the energy function is computationally a very costly operation due to the very large number of interaction terms.

B. The optimization problem

Assume an initial center of rotation p is selected. (We describe a number of options that we have considered in our experiments in Section V).

Let $\mathbf{g} = (R, t) \in SO(3) \times \mathbb{R}^3$. Then, the rigid transformation of $\mathbb{R}^3$ given by $\pi_1 \circ \mathbf{g}_p$ extends naturally to a rigid transformation of the ligand, namely

$$\mathbf{g}_p: \mathbb{R}^{3 \times m_l} \rightarrow \mathbb{R}^{3 \times m_l}.$$

Let $\mathbf{q} = (q_1, \dots, q_{m_l}) \in \mathbb{R}^{3 \times m_l}$ denote a position of the ligand and $\mathbf{q}' = (q'_1, \dots, q'_{m_l}) \in \mathbb{R}^{3 \times m_l} = \mathbf{g}_p(\mathbf{q})$. Then, $q'_i \in \mathbb{R}^3$ is defined as

$$q'_i = \pi_1 \circ \mathbf{g}_p(q_i).$$

Therefore, the local optimization problem can be defined as the following optimization on $SO(3) \times \mathbb{R}^3$

$$\min_{\mathbf{g}} E(\mathbf{g}_p(\mathbf{q})) \quad \mathbf{g} \in SO(3) \times \mathbb{R}^3.$$

C. The optimization approach

Our approach to solving the above optimization problem is driven by practical considerations and follows one of the standard options available for manifold optimization, namely using a local parametrization of the manifold. This parametrization is given by the exponential map on the tangent space of $SO(3) \times \mathbb{R}^3$ at $(I, \mathbf{0})$ (see, e.g. [4]). It is wellknown that using local coordinates introduces nonlinearities that can degrade the performance of optimization algorithms. Hence, there is a tradeoff in using local parametrizations (see, e.g., [2]). Our motivation for using a local parametrization is to be able to take full advantage of well-developed and optimized Euclidean optimization algorithms. Our evaluation of the tradeoff is driven by a very practical concern; namely, we aim to develop an algorithm of equal or higher quality that is more efficient than current algorithms used in protein docking. Our experimental results, given in Section V show that our approach has been successful.

One can also justifiably argue that the local parametrization based on the exponential map is particularly well-suited for the local optimization we consider. The fact that our rigid body

optimization is defined on $SO(3) \times \mathbb{R}^3$ rather than $SE(3)$ is another argument in favor of the suitability of exponential parametrization.

D. The Optimization algorithm

Given the exponential map parametrization, the rigid body energy minimization is defined on the 6-dimensional Euclidean space $\mathbb{R}^6$. From among the many deterministic algorithms available to solve local minimization problems on a Euclidean space, we have selected the quasi-Newton method of Limited memory BFGS (LBFGS) [15]. In our parametrization, the gradient and the Hessian of the energy function with respect to the parameters of optimization can be explicitly calculated. However, these are costly operations, evaluating the Hessian being significantly more costly than evaluating the gradient. Our choice of LBFGS has been based on the fact that it uses only gradient information to obtain second order information about the energy function.

Denoting the elements of $\mathbb{R}^6$ by x , the LBFGS method consists of the following iterations[15]

$$x_{k+1} = x_k + \alpha_k d_k, \quad (1)$$

where

$$d_k = -H_k \nabla E_k, \quad (2)$$

where ∇E_k is the gradient of the energy function, H_k is the LBFGS approximation of the inverse of the Hessian of the energy function as described in [15], and α_k is an appropriately selected step-length as described in [16].

To avoid moving away from a local minimum that is in the vicinity of the initial configuration, we avoid big rotational moves in the iterations of the algorithm. In the initial configuration there may be clashes between the ligand and the receptor and the energy and its gradient may be very high; As a result it is possible that at the first step the algorithm may suggest a big rotational move. In such cases, we scale the diagonal elements of the initial Hessian approximation corresponding to the rotational parameters to avoid big rotational moves. At subsequent steps, if the algorithm suggests making a big rotational move, we re-initialize the Hessian to the identity matrix and restart LBFGS.

Figure 1 (a) & (b) give a schematic of our parametrization approach. The local optimization is performed on the tangent space. Figure 1 shows the evolution of the optimization algorithm on the tangent space until a local minimum is reached. The solution is then mapped to the manifold of rigid body transformations. Figure 1 (b) shows the evolution of the optimization algorithm in terms of the movement of the ligand. The ligand is shown by a small sphere with an attached coordinate frame that shows its orientation. Translational moves can be seen by the movement of the center of the sphere and rotational moves by the rotation of the coordinate frame.

Figure 2 presents the configuration of the receptor and ligand for the complex 1AY7 before and after the application of the local minimization.

V. EXPERIMENTAL RESULTS

In this section we describe the experimental setup and results from the application of our optimization algorithm to protein-protein docking and compare the performance of our algorithm with an state of the art algorithm used for the same purpose. Our comparison is based on the quality of the solutions generated and the computational efficiency of the algorithms.

The results reported here are based on the application of our algorithms to 9 enzyme-inhibitor, 6 antigen-antibody, and 4 other complexes selected from the protein docking benchmark set [17]. We compare our algorithm with one of the commonly rigid body minimization algorithms, namely CHARMM rigid body minimization [18].

As discussed earlier, in our algorithm we have the flexibility of selecting a center of rotation for rigid body transformation. We examined two different centers of rotation: (i) the center of mass of the ligand and (ii) the center of mass of the contact residue interface of the ligand. The contact residue interface of the ligand is defined as the residues of the ligand which have at least one atom within 10 Å of an atom of the receptor. Our experiments showed that option (ii) produced better results. These results are reported in what follows.

For the quality of the solutions, we consider the ensemble of 1500 solutions produced for each protein pair. The solutions where the local minima found by the two algorithms are within 0.01 Å RMSD distance of each other, or where the difference between the energies of

the solutions found are less than $0.01 \frac{kcal}{mol}$ are considered as ties. If the local minimum found by one of the algorithms is further than 10 Å from the initial conformation, the solution is considered a failure, as we expect to find some local minimum within a 10 Å RMSD range of the initial conformation. The cases where both algorithms fail and there is no basis for comparison are removed from those reported. In all other cases, the quality of the solution of one algorithm relative to the other is considered as superior if it has a lower energy (by more than $0.01 \frac{kcal}{mol}$). For each complex, the number of cases where one algorithm was found to be superior to the other as well as the number of ties are reported in Table I.

As for the measure of computational efficiency of each algorithm, we have selected the number of energy function evaluations needed to converge to a local minimum. Given that energy function evaluations are the most costly operations, the number of energy evaluations is used as a measure of run time efficiency of the algorithm. Since the same energy function is used for both algorithms, the number of energy function evaluations is a fair comparison of the runtime of the two algorithms.

Results for the comparison of the two algorithms with the center of rotation as the center of mass of the contact residue interface are shown in Table I. Each complex is identified by its 4-letter PDB [19] code in the first column of the table. The second column identifies the type of the complex (E: Enzyme-inhibitor, A: Antigen-antibody, and O: Other). The third column gives the number of conformations in which CHARMM (denoted by CH) converged to a local minimum with lower energy than one produced by our algorithm (denoted by MO).

The fourth column presents the number of cases in which our algorithm was superior to CHARMM and the fifth column gives the number of cases where the two algorithms performed similarly. The sixth column gives the average number of energy function evaluations in CHARMM and finally the last column is the average number of energy function evaluations of our algorithm.

Based on the results reported in Table I it can be seen that our proposed algorithm has a better performance, based on the quality of solution criterion, but more importantly, it is on the average about 7.4 times faster than CHARMM.

VI. CONCLUSIONS

In this paper we introduce a new formulation for rigid body minimization that is based on a new group of rigid body transformations. In this new formulation, the rigid body minimization is naturally defined on the Lie group $SO(3) \times \mathbb{R}^3$ which is a direct product of the Lie groups $SO(3)$ and $\mathbb{R}^3$. A key advantage of this new formulation over the common formulation of rigid body optimization as an optimization on $SE(3)$ is that the complications associated with optimization on $SE(3)$ are fully avoided.

The new formulation is used to develop a local manifold optimization algorithm for molecular docking. Experimental results provided in the paper show that our algorithm is substantially more efficient than a state of the art alternative for protein-protein docking.

We are currently considering the following directions as a continuation of the research reported in this paper: (a) the development of an optimization algorithm that evolves directly on the $SO(3) \times \mathbb{R}^3$ manifold rather than a local parametrization of it, (b) a more rigorous evaluation of the performance of the resulting algorithms, and (c) developing molecular docking algorithms under some molecular flexibility.

References

- [1]. Li Z, Guo J, Chu Y. Geometric algorithms for workpiece localization. *IEEE Transactions on Robotics and Automation*. 1998; 14(no. 6):864–878.
- [2]. Gwak S, Kim J, Park FC. Numerical optimization on the Euclidean group with applications to camera calibration. *IEEE Trans. on Robotics and Automation*. Feb.2003 19(no. 1):65–74.
- [3]. Tron, R.; Vidal, R. Distributed image-based 3-d localization of camera sensor networks. *Proceedings of the 48th IEEE Conference on Decision and Control*; 2009. p. 901-908.
- [4]. Selig, JM. *Geometric Fundamentals of Robotics*. Springer; 2005.
- [5]. Smith ST. Optimization techniques on Riemannian manifolds. *Proc. Fields Inst. Workshop on Hamiltonian and Gradient Flows, Algorithms, and Control*. 1994
- [6]. Absil, PA.; Mahony, R.; Sepulchre, R. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press; Princeton, NJ: 2008.
- [7]. Murray, RM.; Li, Z.; Sastry, SS. *A Mathematical Introduction to Robotic Manipulation*. CRC Press; 1994.
- [8]. Chirikjian GS. Group theory and biomolecular conformation: I. Mathematical and computational models. *Journal of Physics: Condensed Matter*. 2010; 22(no. 32)
- [9]. Halperin I, Ma B, Wolfson H, Nussinov R. Principles of docking: An overview of search algorithms and a guide to scoring functions. *Proteins*. 2002; 47:409–443. [PubMed: 12001221]

- [10]. Smith G, Sternberg M. Prediction of protein-protein interactions by docking methods. *Curr Opin Struct Biol.* 2002; 12:28–35. [PubMed: 11839486]
- [11]. Vajda S, Kozakov D. Convergence and combination of methods in protein-protein docking. *Curr. Opin. Struct. Biol.* Apr.2009 19:164–170. [PubMed: 19327983]
- [12]. Paschalidis IC, Shen Y, Vajda S, Vakili P. Sdu: A semi-definite programming-based underestimation method for stochastic global optimization in protein docking. *IEEE Transaction on Automatic Control.* 2007; 52(no. 4):664–676.
- [13]. Shen, Y.; Vakili, P.; Vajda, S.; Paschalidis, IC. Optimizing noisy funnel-like functions on the Euclidean group with applications to protein docking. *Decision and Control, 2007 46th IEEE Conference on*; 2007. p. 4545-4550.
- [14]. Shen Y, Paschalidis IC, Vakili P, Vajda S. Protein docking by the underestimation of free energy funnels in the space of encounter complexes. *PLoS Computational Biology.* 2008; 4(no. 10): 843865.
- [15]. Liu DC, Nocedal J. On the limited memory BFGS method for large scale optimization. *Mathematical Programming.* 1989; 45:503–528.
- [16]. Xie D, Schlick T. A more lenient stopping rule for line search algorithms. *Optimization Methods and Software.* 2002; 17:683–700.
- [17]. Hwang H, Pierce B, Mintseris J, Janin J, Wang Z. Protein-protein docking benchmark, version 3.0. *Proteins.* 2008; 73(no. 3):705–709. [PubMed: 18491384]
- [18]. Brooks BR, Brucoleri RE, Olafson DJ, States D, Swaminathan S, Karplus M. CHARMM: A program for macromolecular energy, minimization, and dynamics calculations. *Journal of Computational Chemistry.* 1983; 4:187–217.
- [19]. Berman H, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The protein data bank. *Nucleic Acids Research.* 2000; 28:235242.

Fig. 1.

(a) The sphere represents the $SO(3) \times \mathbb{R}^3$ manifold and the plane represents the tangent space at the identity. The dots on the tangent space correspond to optimization steps and the position of each dot corresponds to the first two coordinates of the exponential map parametrization at the identity. The position of the local optimization algorithm on the tangent space after every ten steps is shown by a color dot. Colors correspond to the energy value at that step of the optimization. Red represents high energy and blue represents low energy. Each step of the optimization is connected by a line to the next step. (b) Each sphere represents the center of mass of the ligand at every ten step of the optimization of the 1AY7 complex. The color codes are the same as in (a). The axes connected to each sphere show the rotational axes of the ligand at that step of the optimization.

Fig. 2.

- a) 1AY7 complex before rigid body minimization, the coordinate axes is centered at the center of rotation. b) 1AY7 complex after rigid body minimization, the axes rotate and translate with the ligand and sit in a new place

TABLE I

Comparison of the quality of solutions & computational efficiency of our manifold optimization (MO) WITH CHARMM RIGID BODY MINIMIZATION (CH)

Complex description		Quality of solutions: Which performs better			Computational efficiency: Average no. of steps	
Complex	Type	$CH > MO$	$CH < MO$	$CH = MO$	CH	MO
1AVX	E	95	240	890	1027	111
1AY7	E	259	267	374	1650	116
1EAW	E	177	276	900	578	93
1MAH	E	316	319	364	869	134
1PPE	E	202	246	1044	453	125
1R0R	E	154	284	1020	638	113
2PCC	E	328	377	472	990	143
2SIC	E	351	164	479	754	104
2SNI	E	281	222	696	599	110
1FSK	A	173	475	814	1145	110
1NCA	A	85	182	875	1726	121
1WEJ	A	344	443	630	820	125
2JEL	A	432	328	507	946	120
1E6J	A	228	260	893	606	102
1AHW	A	188	595	489	1285	105
1B6C	O	165	838	230	305	102
1BUH	O	179	688	503	209	117
1GLA	O	114	880	276	227	98
1GPW	O	90	954	252	1094	101
		17.4%	33.6%	49.0%	7.4	1