

A Method for Biomolecular Structural Recognition and Docking Allowing Conformational Flexibility

BILHA SANDAK,¹ RUTH NUSSINOV,^{2,3} and HAIM J. WOLFSON⁴

ABSTRACT

In this work, we present an algorithm developed to handle biomolecular structural recognition problems, as part of an interdisciplinary research endeavor of the Computer Vision and Molecular Biology fields. A key problem in rational drug design and in biomolecular structural recognition is the generation of binding modes between two molecules, also known as molecular docking. Geometrical fitness is a necessary condition for molecular interaction. Hence, docking a ligand (e.g., a drug molecule or a protein molecule), to a protein receptor (e.g., enzyme), involves recognition of molecular surfaces. Conformational transitions by "hinge-bending" involves rotational movements of relatively rigid parts with respect to each other. The generation of docked binding modes between two associating molecules depends on their three dimensional structures (3-D) and their conformational flexibility. In comparison to the particular case of rigid-body docking, the computational difficulty grows considerably when taking into account the additional degrees of freedom intrinsic to the *flexible molecular docking* problem. Previous docking techniques have enabled hinge movements only within small ligands. Partial flexibility in the receptor molecule is enabled by a few techniques. Hinge-bending motions of protein receptors domains are not addressed by these methods, although these types of transitions are significant, e.g., in enzymes activity. Our approach allows hinge induced motions to exist in either the receptor or the ligand molecules of diverse sizes. We allow domains/subdomains/group of atoms movements in either of the associating molecules. We achieve this by adapting a technique developed in Computer Vision and Robotics for the efficient recognition of partially occluded articulated objects. These types of objects consist of rigid parts which are connected by rotary joints (hinges). Our method is based on an extension and generalization of the Hough transform and the Geometric Hashing paradigms for rigid object recognition. We show experimental results obtained by the successful application of the algorithm to cases of bound and unbound molecular complexes, yielding fast matching times. While the "correct" molecular conformations of the known complexes are obtained with small RMS distances, additional, predictive good-fitting binding modes are generated as well. We conclude by discussing the algorithm's implications and extensions, as well as its application to investigations of protein structures in Molecular Biology and recognition problems in Computer Vision.

¹Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel.

²Laboratory of Experimental and Computational Biology, SAIC, NCI-FCRDC, Fredrick, Maryland.

³Sackler Institute of Molecular Medicine, Faculty of Medicine, Tel Aviv University, Tel Aviv, Israel.

⁴Computer Science Department, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel.

Key words: flexible molecular docking and molecular recognition, induced fit, structure-based drug design, object recognition, Hough transform, Geometric Hashing.

1. INTRODUCTION

WE PRESENT AN ALGORITHM TO PREDICT INTERACTIONS between biomolecules of known three-dimensional (3-D) structure with internal degrees of freedom. Specifically, given the 3-D shapes of two biomolecules, we would like to position them in such a way that an interaction occurs. Such an interaction is called "docking." Molecular docking is an essential step in many biochemical activities such as regulatory mechanisms, chemotherapy, and toxicity. The ability to automatically predict molecular binding modes is important, for example, in structure-based drug design. An efficient algorithm for docking prediction can aid in database searches for drugs fitting target enzymes, inhibiting, or altering the enzymes' undesired activity. Our approach has been inspired by previous work in Computer Vision dealing with articulated object recognition (Wolfson, 1991). Although coming from entirely different application fields, the underlying geometric principles bear a similar flavor.

Successful docking requires shape complementarity of the interacting molecules. Thus, usually the molecular docking problem is first approached by determining geometrically favorable solutions. Subsequently, the proposed bound complexes can be checked for their biochemical feasibility. During the process of molecular association, either of the participating molecules, i.e., the ligand or the protein receptor, may undergo conformational changes, enabling their access and binding. The induced conformational transitions involve flexible movements of molecular parts, in the form of rotational movements of relatively rigid subparts about hinges, namely, *hinge-bending* movements (Figure 1). Hinge-induced movements in receptors are important, for example, in enabling access and better fitting of substrates or inhibitors to the active site of an enzyme receptor. Movements of enzyme domains may thus induce catalytic or inhibition processes upon binding substrates or inhibitors (Figure 1C, D). Molecular docking depends on the spatial structures of the associating molecules and on their conformational flexibility. This problem is referred to as the *flexible molecular docking* problem.

Rigid docking is a particular case of the general flexible docking problem, where a rigid ligand ("key"), is docked onto a rigid protein receptor ("lock"). Heuristic rigid body methods have been developed (see, for




FIG. 1. A schematic illustration of hinge-bending. (A) A flexible ligand consisting of three rigid parts connected by two hinges; two consecutive parts share a common hinge. (B) The flexible ligand displayed in A is docked to a receptor (bold-faced line type). (C) A flexible receptor consisting of two rigid domain parts connected by a hinge—"open form." (D) The flexible receptor displayed in C binds a ligand (light line type)—"closed form."

example, Katchalski-Katzir *et al.*, 1992; Jiang and Kim, 1991; Wang, 1991; Kuntz *et al.*, 1982; Fischer *et al.*, 1995). Yet, these methods do not take into account the conformational transitions that molecules may undergo. This short-coming is exemplified by a database search for inhibitors to the HIV-1 protease (DesJarlais *et al.*, 1990), which enables the prediction of only rigid docked conformations. By allowing flexibility in either the ligand or receptor molecules, additional candidate inhibitors may be obtained. These are most likely to be ignored during the limited and restrictive rigid-body docking search, resulting in missing potential therapeutic agents. It is therefore important and beneficial to account for alterations in ligands and receptors molecular conformations when addressing the docking problem.

Previous flexible docking techniques account for induced hinge flexibility only within small ligands and do not enable subparts (domain) motions in receptors (DesJarlais *et al.*, 1986; Leach and Kuntz, 1992; Mizutani *et al.*, 1994; Leach, 1994; Smellie *et al.*, 1991; Goodsell and Olson, 1990; Ghose and Crippen, 1985; Jones *et al.*, 1995; Clark and Ajay, 1995; Jones *et al.*, 1997; Rarey *et al.*, 1996; Welch *et al.*, 1996). Partial receptor flexibility is allowed by Leach (1994), which enables side-chain¹ flexibility, and Jones *et al.* (1995), which allow partial flexibility in hydrogen-bonding groups. None of these techniques enable hinge-induced domain within receptors, although theoretically it seems feasible. Both DesJarlais *et al.* (1986) and Leach and Kuntz (1992) use the DOCK-based approach (Kuntz *et al.*, 1982) for docking a flexible ligand onto a rigid receptor. The ligand is represented as a composition of its rigid parts. Each part is matched separately within the docking site. DesJarlais *et al.* (1986) do not assume dependency on the order in which the ligand's parts are being recognized. However, Leach and Kuntz (1992) have the "most rigid" informative part of the ligand docked in the initial step of the algorithm. The remaining parts are docked using systematic exploration of the ligand space. The method accounts for both shape complementarity and hydrogen-bonding complementarity. Leach (1994) employs a combination of two search methods used for determining the optimal arrangement of amino acid side chains and ligand conformations. The ligand's conformational space is systematically searched. Ghose and Crippen (1985) present a method for determining feasible binding modes of a flexible ligand in the receptor's site using distance geometry. The method superimposes the two bodies, without explicitly calculating the corresponding transformation, according to the intraatomic and interatomic distances constraints. A combinatorial search of the conformational space is carried out. Smellie *et al.* (1991) consider atoms capable of hydrogen bonding. The method is based on clique finding, which is a NP-complete problem. Combinatorial matching is conducted by Mizutani *et al.* (1994) for hydrogen bonds between the groups of a prechosen ligand part and of those of the receptor's. The remaining ligand subparts are examined systematically, reconstructing minimal energy conformations. Goodsell and Olson (1990) employ the Metropolis (simulated annealing) for conformation searching combined with energy evaluations. Since reproducibility and convergence are not guaranteed, a global solution cannot always be obtained. The same difficulty exists in the method of Jones *et al.* (1995). The authors adapt a genetic algorithm that uses a stochastic evolutionary strategy for exploring the full conformational flexibility of the ligand, with partial flexibility of hydrogen bonding groups of the protein receptor. Additional genetic algorithms based approaches are, for example of Clark and Ajay (1995) and Jones *et al.* (1997). Recently, Rarey *et al.* (1996) developed an efficient method (FlexX) for docking flexible organic ligands into protein receptors, by combining a model of the physicochemical attributes of the docked molecules with the sampling of the conformational space. The incremental strategy employed is analogous to the aforementioned strategy of Leach and Kuntz. The docking tool places an interactively chosen base fragment, whereas the additional ligand fragments are incrementally placed into the active site. A similar approach is of Welch *et al.* (1996), which automatically select the base fragment and the docking site.

In this work we describe and apply our general algorithm, which allows conformational flexibility in either the ligand or the receptor molecules. We enable domains/subdomains/group of atoms movements in either of the molecules. Regardless of the existence and the extent of the a priori unknown hinge bending movement and the docking site, our algorithm incorporates the rigid subpart matching technique and the global consistency checks as an integral part of the recognition process. We simultaneously match all parts of the molecule and do not apply a systematic exploration of the conformational space. By considering the molecules as 3-D structures specified by their molecular surface representation, techniques originating in Computer Vision and Robotics can be applied to discover the docked solutions (conformations). The recognition process we face in the docking problem is reminiscent of the automatic part assembly problem in Robotics and of the partially occluded 3-D object recognition task in Computer Vision. A major task in

¹A protein molecule is built up from a repertoire of 20 amino acids. The *side chain* (attached to the central backbone carbon atom) distinguishes between the different amino acids.

Computer Vision is the model-based object recognition problem, which can be formulated as follows: *Given a database of previously observed objects (models), and a newly observed scene (target) with numerous cluttered objects, recover all the occurrences of the database objects in the scene, even if they are partially occluded.* Since in docking we are matching complementary molecular surfaces, the analogy is quite obvious. The database of objects becomes a database of ligands (e.g., drug or protein molecules), and the newly observed scene is the protein receptor. Partial occlusion and additional object clutter is analogous to the fact that only part of the ligand molecular surface binds to a part of the receptor surface with *no a priori* knowledge of the section (patches) that will exhibit the match, i.e., the location of the binding site. The close analogy between rigid docking and object recognition extends also to the flexible case, which is usually referred to in Computer Vision and Robotics as *articulated object recognition*. Articulated objects are objects consisting of rigid parts which are connected either by rotary or sliding (prismatic) joints. The analog of a hinge in the molecule is a rotary joint in an articulated object. The conformational transitions we allow is rotational movements of the molecular substructures about hinges, in either the ligand or the receptor molecules. The extended and adapted recognition task is thus formulated as: *Given a database of known ligands (models), and a newly introduced receptor (target), recover all ligands which exhibit substantial partial surface match with the receptor surface, without colliding with the receptor. If the ligands contain hinges, solve the surface matching problem by recovering the ligand in a plausible conformation, without having the parts self-collide.* The roles of the ligands and the receptors can be interchanged, since the mathematical problem is symmetrical.

Our docking method has originated from previous work in Computer Vision by (Wolfson, 1991), which extended the rigid body matching technique based on Geometric Hashing (Lamdan *et al.*, 1990) and the generalized Hough transform (Ballard, 1981), to an articulated object recognition technique. An implementation of this approach for 2-D application of industrial tool recognition in photographs (such as scissors and pliers) has been reported (Beinglass and Wolfson, 1991). We have successfully applied our adapted 3-D algorithm to interesting molecular complexes cases, in which hinge-bending flexibility is allowed in either the ligand or the receptor molecules, yielding fast matching times of their surfaces. We verify our method by docking molecular complexes of predetermined bound structures.² The correct docked conformations have been obtained, consistent with experimental observations. Furthermore, the transformations which generate the geometrically correct docked solutions, are among the high scoring ones. Predictive, good-fitting, alternate binding modes have been generated as well, for both bound and unbound molecular complexes.

2. THE ALGORITHM

The algorithm allows protein domain motions (swiveling) at hinge regions of the molecule, and hinge rotational movements of ligand subparts. The geometrical acceptable docking solutions are optionally filtered to yield solutions restricted to hinge movements of one degree of freedom around known axes, namely rotational bond movements. Thus, one can choose the general geometric model of a full 3-D rotation, which allows a reasonable approximation of hinge regions of large molecular domains and of few consecutive bonds (due to the short bond length). Alternatively, the restricted geometrical model can be chosen, where the hinge movements are about single rotational bonds.

2.1. Overview

As molecules interact at their surfaces, the ligand and the receptor molecules are described by sets of "interest points," which represent their molecular surfaces. These point sets are matched to achieve sufficient complementarity. The determination of the molecular surface representation employed, i.e., the choice of interest point sets, is a nontrivial task. Two major issues are considered for a discrete point description of the molecular surface. On the one hand, we require the point representation to be as precise as possible, so as to have a high resolution of the surface features. On the other hand, too dense a representation can lead to intractable execution times and high memory consumption, when processing huge point sets. So far, we have used two types of interest point sets, which meet the requirements of being both accurate and sparse. The surface representations we employ have been devised by Kuntz *et al.* (1982), and by Lin *et al.* (1994) and Lin

²A *bound* structure is a structure where the ligand and the receptor have been determined by examining their complex. In case the ligand and receptor molecular structures have been determined separately, it is called an *unbound* case (structure).

and Nussinov (1996). The former describes the surface invaginations and cavities, which facilitates docking flexible ligands onto the active site of the receptor. The latter representation describes key features of the surface by computing critical points for convex, concave and saddle regions of the surface. This description facilitates ligand docking onto hinge bent domain parts of the receptor molecule. It can also describe molecules having relatively "flat" surfaces. Both surface descriptions originate from the solvent accessible surface generated by Connolly (1983a, b) as defined by Richards (1977). This surface is created by rolling a probe ball on the van der Waals spheres of the molecule which produces sets of points, the reentrant (concave) and contact (convex) points. Kuntz generates a set of spheres representing the complementary (negative) from the reentrant and contact points. Each sphere touches the molecular surface at two points and has its center on the surface normal of the first point. The number of spheres are reduced such that only one sphere is retained per surface atom. Overlapping spheres are clustered and the largest cluster of sphere centers represents the receptor's interest point set. The ligand's interest point set are its atom centers. Lin's computes the critical points by projecting the gravity center of points in convex, concave and saddle regions onto the surface, obtaining "cap," "pit," and "belt" points, respectively. Here, cap points comprise the receptor's interest point set and pits are the ligand's.

Our algorithm is divided to two phases, the **preprocessing phase** and the **recognition phase**. An overview of these phases is given below, followed by a detailed description of each phase and its stages. A general outline of the algorithm is depicted in Figure 2. A summary of the application parameters is displayed in Table 1. In order to simplify the exposition, we describe our algorithm for the case where we position the hinge(s) in the ligands, although one can interchange the roles of the ligands and the receptors.

Preprocessing. The ligands information is coded into a look-up (hash) table that serves as the database of ligands. The stored information is invariant to rotations and translations, since the ligands may undergo this type

TABLE 1. PROGRAM PARAMETERS

<i>Parameter name</i>	<i>Definition</i>	<i>Default</i>
<i>collision_distance</i>	A distance criterion for determining whether the van der Waals spheres of the molecules collide	1.75 Å
<i>contact_distance</i>	A distance criterion for determining whether the van der Waals spheres of the molecules are in contact	1 Å
<i>voting_threshold</i>	A percentage criterion for determining which transformations pass the matching-stage and are checked in the collision_check	20%
<i>contact_percentage</i>	A measure for the "goodness" of the binding modes	Dynamically calculated
<i>contact_threshold</i>	A percentage criterion for determining which solutions will be checked in the self_collision_check, depending on their contact_percentage values	80%
<i>run_type</i>	Considering all triplets of receptor's interest points (regular), or segmentized interest point space (rapid)	Rapid
<i>surface_type</i>	Interest point determination—Kuntz's or Lin's molecular surfaces	Lin's
<i>atom_jump</i>	For creating sparseness in the number of atoms checked in the self_collision_check	1
<i>trig_max_len</i>	Maximum distance constraint of the triangle edge length of the triplets of interest point	11 Å
<i>trig_min_len</i>	Minimum distance constraint of the triangle edge length of the triplets of interest point	2 Å
<i>prm_space_reso</i>	The resolution of the voted for parameter space—its bins size	1 Å
<i>hash_tab_reso</i>	The resolution of the look-up (hash) table—its bins size	0.5 Å
<i>verification_type</i>	Considering all high scoring transformations (prune), or clustering them according to their rotations (no-prune)	Prune
<i>verf_cluster_size</i>	The cluster size of the rotations	10°
<i>angle_check</i>	Restrict binding modes to account for hinge-induced rotational bonds movements	Yes
<i>angle_tolerance</i>	Tolerance of the angle spanned by the bonds connected to the hinge	20°




FIG. 2. A general outline of the algorithm. The "model" entity represents the molecular structure of a ligand (or a receptor). The "target" entity represents a receptor (or a ligand) onto which the model(s) are docked. The preprocessing phase can be carried out offline for each new model stored.

of transformation in order to dock to the given target receptor. The hinge locations for every stored molecule are also determined in this phase by employing chemical flexibility considerations. When enabling flexibility in the receptor, the positioning of hinge(s) is determined via comparison of similar structures in different, i.e., open and closed conformations and studying potential hinge regions. As this phase is independent of the receptor molecules, it can be executed off-line.

Recognition. The input to this phase is the receptor surface information. Ligands surfaces having partial fit to the receptor surfaces are recovered, yielding the required transformations for the docking of the ligands' parts onto the receptor. This is achieved by casting a vote for every match between a ligand's surface patch and a receptor's surface patch (*matching stage*). The surface patch is a geometric configuration of the interest points. The hinge position within the docked site is derived from the transformation between the corresponding ligand subpart and the receptor surface patches. The vote is cast, therefore, for the ligand together with the new location of the hinge. The high-scoring transformations (hinge locations) are verified to ensure obtaining geometrically acceptable solutions (*verification stage*). This is done by rejecting the transformations that cause the ligands parts to collide with the receptor structure (collision check) and with each other (*self-collision check*).

We next detail both phases.

2.2. Preprocessing

(a) The ligand is represented as a set of "interest points."

(b) The (known) hinge positions are picked as the origin of 3-D Cartesian coordinate frames, which will be called the "ligand frames." The orientation of these frames is set arbitrarily. The angle spanned by the bonds connected to the hinge is computed for the use of bond rotation constraint, optionally employed in the verification stage.

(c) For each nonordered, noncollinear triplet of interest points, in each ligand part, we define three unique triplet based Cartesian frames, one for each triplet point. When considering a triplet of points, an internal cyclic order is defined, so that the frames are positioned as follows: The origin of each frame is defined at the respective triplet point, the x axis as the line from the point to the neighboring point, the z axis as the normal to the triangle plane obtained by the cross product of the x axis with the second triangle side emanating from the origin, and the y axis as the cross product of the x and z axes. These are the "triplet frames." The distinctive geometric shape signature of the triplet, namely the nonordered triplet of the triangle side lengths serves as an address to a look-up (hash) table. The information stored at this entry is the ligand identification, part number, and the transformations between the "triplet frames" and the "ligand frame." Thus, for example, on a part with two hinges, two sets of transformations will be stored for each triplet point, while on a part with one hinge only, one set of transformations will be stored.

A triplet of points defines the shape signature, or the *invariant*. A minimal and maximal distance constraint (*trig_max.len* and *trig_min.len*, respectively) are introduced. The maximal distance constraint serves to reduce the number of triplets stored and matched, and the minimal distance constraint serves to improve the numerical

robustness of the reference frame computation. A triplet is considered for matching, only if the distance between the composing interest points are within given constraints. This proximity heuristic is based on an assumption that a match between two objects contain relatively dense regions of matching points, and is not composed only of isolated points distant in space. The values of the distance constraints should be large enough to provide for reliable matching. Fewer look-up table records are allocated when posing the distance constraints, due to the reduction in the number of triplets (invariants) considered. Hence, fewer table records are processed in the sequel.

2.3. Recognition

(a) The molecular structure of the receptor is represented by its set of "interest points."

(b) All ordered noncollinear triplets of the "interest points" are considered for the matching stage. For each such triplet, the triplet based Cartesian frames are calculated as explained above. The lengths of the triangle sides are computed. Note that this calculation is invariant under rotation and translation, thus (almost) congruent ligand triangles should have very close values. The look-up table is addressed according to the computed triplet of distances. For each ligand-record present at that entry in the table, the "candidate ligand frames" are computed by applying the prerecorded transformations at that entry to the receptor "triplet frame." The origins of the "candidate ligand frames" are the candidate *hinge locations*. Votes are cast for the identity of the ligand molecule together with the locations (and orientations) of the "candidate ligand frames."

There is an option to have all triplets of receptor interest points considered (*run_type* = "regular run"), or apply pruning, such that the number of triplets are drastically reduced (*run_type* = "rapid run"). This is achieved by dividing the receptor's interest point set to eight segments (octants), and an adjoining ninth segment, whose "corners" are positioned at the geometrical center of each octant. The matching is then conducted for the triplets within each segment. The adjoining ninth segment, which partially overlaps all eight octants, comprises triplets from two or more octants. When constructing the triplets of the ninth segment, triplets not shared by more than one octant are discarded to avoid ambiguity with the already processed triplets of the considered octant. As in the preprocessing phase, the triplets distance constraints are introduced as well.

(c) Finally, the accumulator of votes is searched for high-scoring pairs of (*ligand, hinge location*). Since we are interested only in high-scoring transformations, we pick hinge locations receiving a large number of votes from both parts connected by it. The hinge location is actually the 3-D translation from the original hinge position of the ligand to its new candidate location. The high-scoring hinge locations are determined according to the *voting_threshold* which is a minimal percentage value of the number of votes received by the highest scoring hinge location. Clustering the hinge locations is achieved by the discretization of the continuous 3-D parameter space of the respective translations. The size of the 3-D bins of the parameter space is determined according to the *prm_space_reso* parameter.

(d) A potential match implies existence of complementarity between receptor-ligand surface patches. However, other regions of the two molecules may collide. In the verification stage, the respective transformations of each of the parts, between the initial "ligand frame" and the computed "candidate ligand frame," are applied to the atoms in each of the parts of the ligand. Transformations which result in the penetration of a ligand part into the receptor (collision check) or yielding collisions between the parts of the ligand (self-collision check), are discarded.

In order to speed up the collision check, the ligand part transformations are optionally clustered according to their rotations for every candidate translation. We quantify the degree of rotation using the *angular distance* value of the rotation matrix. This value corresponds to the rotation angle around an equivalent axis, computed as $\arccos \frac{trR-1}{2}$, where R is the rotation matrix of the transformation, and tr is its trace. The rotations having the lowest angular distance in every cluster, are chosen as the representative rotations of the corresponding clusters. These rotations and their respective translations are applied to the ligands parts checked for collision with the target receptor. The parameter which specifies the angular distance cluster size is *verf_cluster_size*. It is the user's choice whether to have the collision check carried out for all transformations (*verification_type* = "no-prune"), or to apply pruning of the rotations via the clustering mechanism (*verification_type* = "prune").

The receptor and the ligand molecules are assumed to collide, if the distance between a ligand atom and a receptor atom is smaller than the sum of their respective van der Waals radii minus a proximity threshold (*collision_distance*). The same criterion is applied to the ligand parts. The van der Waals radii values have been adopted from Kuntz *et al.* (1982) and are presented in Table 2. To speed up the collision check, we reduce the

TABLE 2. VAN DER WAALS RADII

Atom name	Radius
Hydrogen	1.4 Å
Carbon	1.8 Å
Nitrogen	1.8 Å
Oxygen	1.5 Å
Phosphorus	1.8 Å
Sulfur	1.8 Å

size of the space checked. The receptor molecule is divided into eight segments (octants) sharing the geometric center of the molecule. The collision check between a ligand atom and a receptor atom is conducted only in the appropriate receptor's segment (octant), which consists of the respective receptor's atoms.

The "goodness" of a solution is evaluated by employing a score which is based on the number of the ligand's van der Waals spheres which are in contact with the receptor spheres. We refer to this score as the *contact_percentage*. A ligand sphere is assumed to be in contact with a receptor sphere if the distance between the ligand atom and a receptor atom is smaller than the sum of their respective van der Waals radii plus a proximity threshold (*contact_distance*).

Only binding modes receiving a *contact_percentage* that is higher than the *contact_threshold*, are considered for the self-penetration check. These high contact binding modes are optionally passed to a "chemical check" verifying whether the allowed hinge bending motions are restricted to rotational bond movements. This is done by computing the angle(s) spanned by the bonds connected to the hinge (s) in the current binding mode handled. If the angle(s) are similar to the angle(s) computed for the stored ligand in the look-up table (see Preprocessing stage), the binding mode is passed to the self-collision check. Else, it is discarded.

When testing the acceptable binding modes for self penetration, there is an option to reduce the number of atoms checked by using the *atom_jump* constraint. It specifies the relative number of the next atom to be checked. For example, if *atom_jump* equal 3, it means that every third atom in the molecule's part is checked. This parameter is important when conducting the *self-collision check* for very large molecules, such as enzyme receptors. Otherwise, for a small molecule, the parameter has no effect (i.e., its value is set to one). As mentioned above, the self-collision check employs the same criterion for rejecting self-penetration causing transformations, as being done by the collision check.

2.4. Complexity analysis

Since the roles of the ligands and the receptors can be interchanged, the molecular structures stored in the database (look-up table) are referred to as "models." The molecular structure introduced in the recognition phase, onto which the models are docked, is referred to as the "target."

Our research concentrates on the case where the look-up table contains only one model. The complexity described in this section is analyzed accordingly. The notations used throughout this description are summarized in the nomenclature of Table 3.

2.4.1. Preprocessing phase complexity. The shape signature of the object or the object invariant, is defined as a triplet of interest points, allowing a unique definition of a 3-D rotation and translation that superimposes a model structure onto the target structure. We consider nonordered triplets.

Assuming a model consisting of N equal parts, I_m , the number of invariants obtained is

$$I_m = N \times \frac{m/N(m/N - 1)(m/N - 2)}{6}, \quad (1)$$

where m is the number of interest points in the model. The number of invariants is reduced by $1/N^2$ in comparison to the rigid model. The complexity of the term I_m , can be written as

$$I_m = O\left(\frac{1}{6N^2} \times m^3\right). \quad (2)$$

The complexity of the number of invariants imply that the execution times of the program are reduced as the model under study is divided to more and more parts. Since fewer model invariants are stored in the look-up

TABLE 3. NOMENCLATURE

<i>Symbol</i>	<i>Meaning</i>
m	Number of model interest points
n	Number of target interest points
m_A	Number of model atoms
n_A	Number of target atoms
N	Number of models parts
L	Number of models
I_m	Number of model invariants
I_t	Number of target invariants
I_{tF}	Number of target invariants in a regular run
I_{tR}	Number of target invariants in a rapid run
C_p	Complexity of the preprocessing phase
C_m	Complexity of the matching stage
C_{cc}	Complexity of the collision check
C_{scc}	Complexity of the self collision check
B	Number of bins in the look-up (hash) table
D	The difference between the distance constraints
q	The look-up table resolution
R	The average number of records in a look-up table entry bin
f	Number of candidate part transformations passing the voting threshold criterion
f'	Number of candidate part transformations passing the collision check
g	Number of candidate binding modes passing the contact threshold criterion and the optional rotational bonds restriction

table in the preprocessing phase, less model invariant records are handled in the subsequent recognition phase; thus, fewer candidate transformations are considered in the process.

For a small number of hinges, $1/6N^2$ is constant; we obtain that

$$I_m = O(m^3). \quad (3)$$

The insertion time of a model invariant (a record) to the look-up table is assumed to be $O(1)$. In general, for a database consisting of L models, the complexity of the preprocessing phase is of order $L \times m^3$. For a database consisting of a single model, we obtain that complexity of the preprocessing phase, C_p , is

$$C_p = O(m^3). \quad (4)$$

2.4.2. Look-up (hash) table manipulation complexity. The models' information is stored in the look-up (hash) table, so that the models description is memorized for the recognition phase. The data structure considered for the look-up table is such, that a short access time to the table and its entries is enabled. Although the preprocessing phase can be executed off-line, we require the insertion time of a model invariant record in the appropriate entry, to be as rapid as possible. Moreover, in the recognition phase, we require each target (scene) invariant "query" to find its corresponding model invariants in the shortest time possible, since each entry in the table may contain more than one record within it.

Since we have defined three inter-point distances as an address to the look-up table, the table is a three dimensional structure, quantized (or discretized) into bins having a resolution denoted as q . The resolution value is determined by the *hash_tab_reso* parameter. Each bin covers the 3-D interval of q^3 . The number of bins, B , in the table, for a given distance constraints set *trig_max_len* and *trig_min_len* is

$$B = \left(\frac{D}{q}\right)^3, \quad (5)$$

where D is the difference between the maximum distance constraint (*trig_max_len*) and the minimum distance constraint (*trig_min_len*).

The average bin occupancy, i.e., the average number of records stored in each table entry or the average number of invariants stored in a bin, assuming a homogeneous distribution is

$$R = \frac{I_m}{B}, \quad (6)$$

where I_m is the number of model invariants. In case of nonhomogeneous invariant distribution, one can artificially limit the size of the maximal bin by ignoring bins whose size is above a certain threshold. This is justified by the fact that very large bins represent invariants that are not salient enough to contribute significant information to the matching process.

The insertion time of a record to an entry in the table is of order 1. A record is entered as the first record, even if many records already reside in the entry. The average number of operations carried out for each target (scene) query, i.e., the number of operations required for accessing the records stored in an entry, is R , assuming a homogeneous distribution.

2.4.3. Recognition phase complexity. Since the recognition phase is composed of the matching stage that is the core of our algorithm, and the subsequent verification stage, the complexities of the stages are analyzed separately.

Matching stage complexity. The matching stage considers all interest points of the target if *run_type* of regular run is chosen. That is, all ordered triplets of interest points are considered. Hence, the number of target invariants is

$$I_{t_F} = n(n-1)(n-2), \quad (7)$$

where n is the number of interest points of the target. The following complexity approximation is thus written to I_{t_F}

$$I_{t_F} = O(n^3). \quad (8)$$

Another option to process the target structure is to choose a *run_type* of a rapid run, i.e., the target interest point space is divided to eight segments (octants) and an adjoining ninth one. Assuming equal number of interest points per octant, the number of target invariants (ordered triplets) is

$$I_{t_R} = 9 \times n/8(n/8-1)(n/8-2), \quad (9)$$

giving the complexity approximation of

$$I_{t_R} = O\left(\frac{1}{57} \times n^3\right). \quad (10)$$

In general, we can write that the approximation of the complexity of the number of target invariants in the recognition is

$$I_t = O(n^3). \quad (11)$$

Based on Equation (11), we can write that the complexity of the matching stage is

$$C_M = O(n^3 \times R), \quad (12)$$

where R is the average access time to a record in a look-up table entry. The actual complexity thus depends on the size of the bins of this table. Large bins are obviously undesirable, since they do not contribute much to the discrimination process as well. In this case, we “vote” for many candidates simultaneously. If one avoids extremely large bins, this access time will be constant, thus achieving practical matching stage complexity of

$$C_M = O(n^3). \quad (13)$$

However, it should be taken into consideration that there may be a trade-off when defining small bins and consequently reducing the value of R . As the resolution of the look-up table has a dependency on the accuracy of the data stored, small bins may cause the program to lose the correct solution, since in this case, small inaccuracies in the input data may not be tolerated.

Verification stage complexity. The complexity of the verification stage is composed of the complexities of the collision and self-collision checks.

The complexity of the *collision check* is

$$C_{cc} = O(m_A \times n_A \times f(\text{voting_threshold}, I_t, R)), \quad (14)$$

where m_A is the number of docked model atoms, and n_A is the number of target atoms. The term $f(\text{voting_threshold}, I_t, R)$ represents the number of candidate transformations passing the voting threshold criterion. That is, only for these transformations the model part is transformed and checked for collision within the target structure it is docked into. The transformations are deduced from the target invariants and look-up table occupancy, hence the number of transformations depends on the number of target invariants I_t , on the average look-up table occupancy per entry R , and on the voting threshold. Since the molecular space of the target is divided to eight segments, and the model may be divided to N parts (segments and parts assumed equal in this discussion), the complexity is

$$C_{cc} = O\left(\frac{m_A}{N} \times \frac{n_A}{8} \times f\right), \quad (15)$$

thus, speeding up the check.

The number of transformations passing the voting threshold f , is composed of the transformations new hinge locations (translations) and their rotations. By clustering the rotations, the number of transformations checked is significantly reduced as only one representative is chosen from each cluster of the size *verf_cluster_size*. Thus, the order of the number of operations carried in the *collision check* is further reduced.

In the *self-collision check*, the docked model parts are checked for self-penetration. Assuming the model is equally divided to N parts, the average number of operations carried out

$$C_{scc} = \binom{N}{2} \times \frac{m_A^2}{N^2} \times g(\text{contact_threshold}, f'), \quad (16)$$

where the first term in the equation corresponds to the number of times required to apply the self-collision check to every two parts of the molecule. Since in our current application we are handling molecules with two and three parts, this number is very small. The term $g(\text{contact_threshold}, f')$ represents the number of high contact candidate binding modes that have passed the collision check. As such, this number depends on the number of transformations passing the collision check f' , and on the contact threshold criterion which defines a contact percentage value for binding modes to be considered. If one considers only the binding modes which have undergone rotational bond movements, then the number of candidate binding modes checked for self-penetration is reduced. Introduction of the *atom_jump* parameter can drastically reduce the number of operations executed, since by definition every *atom_jump* - 1 atoms are skipped during the check to deliberately cause acceptable sparseness. Thus, the average number of operations is reduced to

$$C_{scc} = \binom{N}{2} \times \frac{1}{\text{atom_jump}^2} \times \frac{m_A^2}{N^2} \times g. \quad (17)$$

The *atom_jump* parameter is not used in the collision check, since, in this process, the contact percentage of the model part is calculated and it is necessary to consider all atoms in the calculation.

We have tested a grid-based approach for conducting the collision check. The target structure is "placed" on a grid, and the transformed model atoms are checked for collision with the target atoms. This is carried out by checking whether the bins corresponding to the van der Waals radii of the query model atoms intersect with the already occupied bins of the target van der Waals spheres. This approach proved inefficient both in memory and time. It is inappropriate for the self-collision check, as the molecular subparts are transformed for every configuration checked. Recently, Halperin and Overmars (1994) have developed an algorithm for conducting intersection queries in molecules. The complexity of their approach depends on the data structures used. The perfect hashing data structure, which describes the target molecule composed of n balls, is constructed using $O(n)$ space and randomized preprocessing time. The intersection queries are answered in time $O(1)$ for balls whose radii are not greater than the radius of the largest ball. This algorithm should be investigated for its efficiency, robustness and memory consumption when employed on our molecular cases, as the molecules we dock have a highly diverse range of sizes.

3. EXPERIMENTAL RESULTS

The atom coordinates considered as input to our algorithm, have been determined by the x-ray crystallography and nuclear magnetic resonance (NMR) techniques, and stored in the Brookhaven Protein Data Bank (PDB) as 3-D molecular structures (Bernstein *et al.*, 1977; see Branden and Tooze, 1991, for a description of the structure determination techniques). In order to verify our algorithm, we have investigated *bound* molecular structures. This type of complexes comprise the 3-D structures (atom coordinates) of ligand(s) and receptor molecules bound together. By applying our method to molecules extracted from bound configurations, we reproduce the binding mode which is in agreement with the empirical observation. Since the ligand and the receptor have been extracted from bound complexes, the "correct" geometrical solutions are those with rotations and translations close to zero. The results we have obtained for the bound cases have small root mean squared distances (*RMSD*), as compared with their native crystal/NMR structures. We define the term "best solution," as the solution having the lowest *RMSD* values of the molecule-parts. This definition serves as a measurement for the performance of our algorithm when verifying it. It has no meaning in solving "real" docking problems, where the input is the ligand and the receptor molecular structures are determined separately. A general docking score is used, having the solutions of the bound and unbound cases ranked according to the size of the contact area(s) of the docked ligand and receptor molecules. It is referred to as the contact percentage [see definition in step (d) of the Recognition-phase of the algorithm] and used to rank the acceptable binding modes. In addition to the correct solutions obtained by our method, good-fitting predictive binding modes have been generated as well. Fast matching (recognition) times have been obtained for all cases studied. The executions of the program have been conducted on the SGI R10000 machine.

Previously, we have docked hinge-bent ligands of various molecular complexes such as the NADPH ligand and the dihydrofolate reductase (DHFR) receptor, using a preliminary implementation of our algorithm (Sandak *et al.*, 1995). In Sandak *et al.* (1996b), we have investigated docked configurations of the U75875 inhibitor and the HIV-1 protease by positioning the hinge at different locations on the inhibitor molecule. In Sandak *et al.* (1996a), we allow hinge induced conformational flexibility in either a peptide ligand or the calmodulin receptor molecule, for both complexed (bound) and unbound receptor structures. As we have extended our method to account for multiple (double) hinges, we also explore this case. Here, we apply our extended method to diverse molecular cases. We summarize the results obtained in Table 4, which appears in the "Performance Summary" subsection (see also Sandak, 1997). Following this section is a detailed description of docking investigations and the biological functions molecules of two cases, which correspond to the MTX/DHFR and MBP/maltose columns of Table 4. The first case (the methotrexate ligand and the DHFR receptor) accounts for docking a flexible ligand onto to a rigid receptor. An unbound docking investigation is conducted for the maltose binding protein receptor and the maltose ligand, where receptor conformational flexibility is permitted. The results of the other eight cases are discussed in brief in the subsection "Additional Cases."

3.1. Performance summary

A summary of the performance measurements appears in Table 4. Following is a description of the table's entries.

Input size. The first four rows of the table correspond to the input size of the different molecular cases explored. The entries comprise of: m , the number of model interest points; n , the number of target interest points; m_A , the number of model atoms and n_A , the number of target atoms.

Molecular cases under study. The first six columns of both parts of the table comprise the cases where the hinge is positioned in the ligand. The seventh column summarizes the case where two hinges are introduced in the ligand. The remaining three columns correspond to the cases where the hinge is positioned in the receptor molecule. Bound (complexed) cases are dealt with in the first eight columns, whereas unbound cases are handled in the last two. The data of bound ligand and receptor are extracted from PDB files describing the structure of the complexed molecules. The data of an unbound pair of molecules are extracted from two different PDB files.

Execution time measurements. The next group of runtime measurements consist of the execution time measured for the different stages of the algorithm. These are the preprocessing phase, the matching stage, the collision check, and the self collision check. The three latter cases comprise the recognition phase of the algorithm. The verification stage in this phase is composed of the collision check and the self-collision check. Although the preprocessing stage can be executed off-line, here the "total" time measurements are the summation of the execution time of the recognition and preprocessing phases.

TABLE 4. PERFORMANCE (RUNTIME MEASUREMENTS) SUMMARY TABLE

<i>Model target</i>	<i>U75875 HIV-1P</i>	<i>NADPH DHFR</i>	<i>MTX DHFR</i>	<i>NAD LDH</i>	<i>Ag Ig</i>
<i>m</i>	59 (36, 24) ^a	48 (27, 22)	33 (13, 21)	51 (23, 29)	58 (31, 28)
<i>n</i>	351	427	427	621	149
<i>m_A</i>	59 (36, 24) ^b	48 (27, 22)	33 (13, 21)	51 (23, 29)	58 (31, 28)
<i>n_A</i>	1517	1294	1294	2560	432
Execution times ^c					
Preprocessing ^d	—	—	—	—	—
Matching stage	0.1	0.1	0.1	0.4	0.1
Collision check	0.32	0.32	0.3	12.13	0.25
Self collision check	—	—	—	—	—
Total	0.42	0.42	0.4	12.43	0.35
Best solution					
RMSD first part	0.62 Å	0.61 Å	1.48 Å	0.83 Å	2.22 Å
RMSD second part	0.96 Å	0.86 Å	2.64 Å	1.33 Å	1.72 Å
Hinge-rank percentage	0.01%	0.25%	0.64%	11.54%	0.94%
Final position	3	1	2	40	5
Total binding modes	5	5	3	767	13

<i>Model target</i>	<i>M13 CaM</i>	<i>M13 CaM</i>	<i>CaM M13</i>	<i>CaM^e M13</i>	<i>MBP^e Maltose</i>
<i>m</i>	210 (103, 108) ^a	210 (87, 37, 92)	526 (265, 261)	465 (227, 238)	804 (330, 474)
<i>n</i>	440	440	284	284	40
<i>m_A</i>	210 (103, 108) ^b	210 (87, 37, 92)	1,164 (572, 593)	1,132 (572, 561)	2,862 (1,244, 1,619)
<i>n_A</i>	1,164	1,164	210	210	23
Execution times ^c					
Preprocessing	0.1	0.08	0.6	0.42	2.06
Matching stage	0.43	0.37	0.35	0.28	0.02
Collision check	1.42	1.55	12.13	38.02	0.37
Self collision check	—	—	—	—	—
Total	2.85	2.00	13.08	38.72	2.45
Best solution					
RMSD first part	2.53 Å	2.03 Å	7.27 Å	4.32 Å	6.45 Å
RMSD second part	1.17 Å	0.98 Å	3.51 Å	5.72 Å	6.23 Å
RMSD third part	—	1.03 Å	—	—	—
Hinge-rank percentage	0.46%	1.51% 2.97%	9%	42%	52%
Final position	2	3	450 (4.5%)	1,562 (0.5%)	107 (9%)
Total binding modes	2	3	10,091	307,982	1,179

^aThe numbers in parenthesis correspond to the number of interest points in each of the model parts.^bThe numbers in parenthesis correspond to the number of atoms in each of the model parts.^cThe times are given in minutes.^dThe notation—means that the phase took under 1 second.^eUnbound cases.


Best solution. The final group of rows accounts for the information relating to the best solution obtained by our algorithm. It is defined as the solution having the smallest RMSD measurement for the parts of the model from its original determined PDB coordinates. The “RMSD” rows depict the RMSD values obtained for the parts of the model. Another measurement of the performance of the algorithm is a matching score based on the number of votes the candidate hinge locations have received. The hinge location of the best solution is ranked according to the number of votes it has received in the matching stage, prior to the verification stage. The “hinge-rank percentage” is computed as the ratio (in percents) between the rank of the hinge location of

the best solution, and the total number of voted for hinge locations. For example, in the U75875 and HIV-1 protease case, the best solution received a voting score which was ranked as the third highest hinge location. Since 20775 hinge locations were voted for, we obtain the "rank percentage" of 0.01% ($3/20775 \times 100$). The "final position" row displays the position of the best solution obtained at the end of the run, ranked according to the contact percentage. This ranking is carried out for the final group of binding modes, which survive the verification stages. These binding modes are the geometrically acceptable solutions, optionally restricted to rotational bond flexibility. Their number appears in the "binding modes" row.

3.2. Docking methotrexate and dihydrofolate reductase

Nucleotides are the building blocks of nucleic acids. The biosynthesis of nucleotides is a vital process, since these compounds are indispensable precursors for the synthesis of both RNA and DNA. Without RNA synthesis, protein synthesis is halted, and unless cells can synthesize DNA, they cannot divide. Nucleotides are also necessary for continuous repair of DNA, which is critical for cell survival. They further play an important role in all major aspects of the metabolism. The inhibitors to nucleotide biosynthesis, are very toxic to cells. The toxicity of the inhibitors has been used to combat cancer, as well as in the treatment of certain diseases resulting from infection by viruses, bacteria or protozoans. *Methotrexate* (MTX) and a number of related compounds inhibit the reduction of dihydrofolate to tetrahydrofolate, a reaction catalyzed by the *dihydrofolate reductase* (DHFR) enzyme. The inhibitor prevents the synthesis of thymidylate in replicating cells. MTX is used as an anticancer drug, preventing the replication of cancerous cells. However, it is highly toxic for all dividing, normal body cells.

We have carried out our docking investigations positioning the hinge at the C9 atom of the MTX ligand (Figure 3). The DHFR receptor is assumed rigid. The atom coordinates of the MTX/DHFR molecular complex have been taken from the 3DFR data entry file of the PDB (Bolin *et al.*, 1982). We have obtained the results depicted by Figure 4. Prior to the verification stage, the total number of candidate hinge locations receiving votes are 5162. The hinge locations are actually the 3-D translations that the part has undergone from the original hinge position to the new ones. These translations are presented as translation distances in the plot. The translation distance is computed as the l_2 -norm of the transformation's translation vector ($\sqrt{x^2 + y^2 + z^2}$). As Figure 4a shows, the correct translations have been obtained among the high scoring ones. This is depicted in Figure 4c, d, where the translation distances are plotted against their voting score. The translation distance of the transformation which yields the lowest RMSD of the MTX parts (i.e., the best solution), received the rank of 33 (Figure 4c) with a voting score of 268 (Figure 4d). This hinge location is thus among the 0.64% highest scoring 5162 candidate hinge locations. The solutions having the highest contact percentages values






FIG. 4. The MTX/DHFR flexible docking results. (a) The number of votes as a function of the translation distance received by the two parts of the MTX ligand. The results displayed here are obtained from the matching stage of the algorithm, i.e., prior to verifying them in the collision and self-collision checks. (b) The "goodness" of the generated binding modes. The percentage of the MTX ligand's van der Waals spheres that are in contact with the DHFR receptor spheres (ordinate) are plotted as a function of translation distances for the different binding modes. (c) The candidate hinge locations represented by their translation distances (ordinate), plotted as a function of their voting score rank (abscissa). (d) The score, i.e., the number of votes received by the corresponding translation distances in c versus the rank. In both c and d, only the first 100 top scoring hinge locations are depicted.

(Figure 4b) correspond to the correct solutions, and to the additional good fitting binding modes. The best solution has been obtained with small RMSD from the original ligand crystal structure as depicted in Figure 5: 1.48 Å for the first part of the MTX ligand and 2.64 Å for the second. The translation distance of the first part is 0.49 Å, and that of the second is 0.36 Å. The respective angular distances are 45.4° and 63.1°. At the end of the run, the best solution is the ranked as the second top scoring solution. Some of the alternate good-fitting binding modes are depicted in Figure 6.

The execution times of the program are 0.1 minutes for the matching stage, 0.3 minutes for the collision check, and under one second for the self-collision check, giving a total of 0.4 minutes. The number of interest points of the DHFR receptor is 427. The combined number of interest points of the two parts of the MTX ligand is 33. There are 13 points in the first part and 21 points in the second. The two parts are joined by the hinge point. The number of atoms in the DHFR receptor is 1294, and in the MTX ligand is 33 (13 in the first part, and 21 in the second part).

3.3. Docking maltose and maltose-binding protein

There are a variety of substrates that are transported through the periplasm. The periplasm is the space between the inner and outer membrane of a bacterium. The substrates transported include peptides, amino acids, vitamins, monosaccharides, and ions. The bacterial periplasmic system, which transport these substrates, consists of initial receptors, the periplasmic substrate-binding proteins, and protein components that translocate the substrate from the periplasm to the cytoplasm. The substrate first binds to the binding protein, and these interact with the membrane-bound complex, which translocates the substrate to the cytoplasm with concomitant ATP hydrolysis (Oh *et al.*, 1993). There are about two dozen periplasmic binding proteins, all of which are monomeric, having two distinct globular domains separated by a deep groove or cleft. Some examples are the LAO (lysine/isoleucine/valine)-binding protein (Oh *et al.*, 1993) and the GBP (galactose/glucose)-binding protein.

Here we study the binding conformations of the *maltose/maltodextrin binding protein (MBP)*. Without the bound maltose, the two domains of the MBP receptor are farther apart and the cleft is wide open ("open form") (Figure 7a). Upon binding the maltose, the ligand is engulfed in the cleft between the domains, such that the ligand-induced conformational change results in a "close form" of the protein receptor. The bound maltose is buried in the groove and is almost completely shielded from the bulk solvent. The two domains, the N-domain and the C-domain, are joined by three linkages, two at the base of the cleft and the third at its side. The N-domain consists of residues³ 1–109 and 264–309. The C-domain comprises residues 114–258 and 316–370. These lobes are connected by the three interdomain segments, which consist of residues 110–113, 259–263, and 310–315, respectively. The main structural distortion upon maltose binding, are mediated entirely by hinge bending of the main chain of the segments. There is a hinge opening about an axis through the first and second segments (residues 111 and 261) (Sharff *et al.*, 1992).

We have studied the binding modes of the unbound MBP receptor to the maltose ligand. The atom coordinates of the unbound MBP have been taken from the IOMP data entry file of the PDB (Sharff *et al.*, 1992). The atom coordinates of the maltodextrin have been extracted from the bound complex of the 2MBP file (Spurlino *et al.*, 1991). We have divided the MBP molecule to two parts according to its domain partitioning. The hinge has been positioned at the C α atom of the Glu¹¹¹ residue, as it is assumed that this residue is significant in the ligand-induced domain motion, located at the base of the binding cleft (Sharff *et al.*, 1992).

In Figure 7b, we depict a potential binding mode we have obtained by our algorithm, of the unbound MBP with the maltose ligand. It is the geometrical solution yielding the highest contact percentage. The maltose molecule is composed of 23 atoms that are engulfed by the MBP receptor. Although the ligand is a very small molecule, the binding mode we portray in Figure 7b, yields a very good fit between the ligand and the receptor. All of the maltose ligand atoms are in contact with the MBP receptor. This result manifests the ability of our method to successfully dock very small molecules having small matching surface, as well as large ones. The unbound MBP has undergone a significant translation and rotation of its two domains, enabling it to transform to a "closed conformation" from its native "open form" (20° difference). It is docked in a reversed fashion as compared with the bound MBP, as the maltose ligand has a somewhat round and symmetrical structure. Owing to their large size, the extent of the hinge-induced movement of the protein domains is not uniform.

³During protein synthesis, the amino acids constructing the molecule are linked to each other as *residues*, creating polypeptide chains.



FIG. 5. The best "correct" docked conformation of the MTX ligand, predicted by our algorithm, is represented as a yellow stick. It is plotted against the pink stick and van der Waals sphere representation of the crystal structure. The DHFR molecule of the crystal complex is drawn as a ribbon. The best solution almost coincides with the crystal structure.

When analyzing the angular motions involved, we observe relatively small movements of the atoms in the vicinity of the hinge, and larger movement farther away. We therefore define the RMSD of interface atoms as the RMS distance of the atoms of one molecule which are in contact with the atoms of the second. The best solution has an average RMSD value of the interface atoms of around 6 Å (results not shown). The running times recorded are 0.02 minutes for the matching stage, 0.37 minutes for the collision check, and under 1 second for the self-collision check, giving the total of 0.39 minutes (excluding the off-line preprocessing stage). The number of interest points describing the maltose ligand is 40. The combined number of interest points composing the two parts of the MBP receptor is 804. The first part consists of 330 points, and the second part of 474 points. The number of atoms in the maltose ligand is 23 and in the MBP receptor is 2,862 (1,244 atoms in the first parts and 1,619 atoms in the second part).

3.4. Additional cases

We have applied our method to five complexes: the HIV-1 protease complexed with the U-75875 inhibitor; the dihydrofolate reductase complexed with methotrexate, and separately with NADPH; lactate dehydrogenase complexed with NAD-lactate; and a Fab fragment of an IgG antibody complexed with a peptide antigen (crystallized as residues 69–87 of myohemerythin). These cases correspond to columns 1–5 in Table 4. The ligands in these cases range from 33 to 59 atoms, and the receptors are built from 432 to 2,560 atoms. For each case, flexible docking was carried out by allowing hinge-bending in the ligand molecules. Applying our method to molecules extracted from bound configurations, we reproduce the binding mode which is in agreement with the experimental observation. Since the ligand and the receptor have been extracted from these types of complexes, the correct geometrical solutions are those with rotations and translations close to zero. Indeed, the binding modes we have obtained have small RMSD, as compared with the native crystal structures. The average RMSD of a correct solution is 1.4 Å, and the average matching time for each complex is around 0.17 minutes. Additional predictive binding modes have been generated as well.

We also apply our approach to the calmodulin (CaM) receptor and its M13 ligand, allowing flexibility in either of the molecules. In all cases, acceptable docked configurations have been achieved. As expected, depending on the location and number of hinges, different movements of the substructural parts are required to attain the optimal surface matching. We analyze the different rotational and translational movements,

bending, rotating, and distorting the backbone. Hence, by using such an approach, we see the many ways in which a given open conformation can close on its respective ligand and *vice versa*, to achieve the optimal matching of the molecular surfaces. First we focus on investigating the binding-modes generated by allowing hinge-bending motions in the M13 ligand, a 26-residue synthetic peptide (441 atoms). NMR, as well as computational studies, have indicated that all peptides usually demonstrate an ensemble of conformations (Theriault *et al.*, 1993). Furthermore, the protein receptor strongly influences the binding conformation of the flexible ligand. We have divided the M13 ligand both to two and to three submolecular parts by positioning either one hinge or two hinges along the molecule. The average RMSD of the correct solution is 1.5 Å and the average matching time is 0.4 minutes (see columns 6 and 7 of Table 4). The CaM molecule is a 148-residue protein (2,259 atoms), involved in cellular Ca^{2+} -dependent signaling pathways. We have explored the binding modes of the bound and unbound CaM receptor to the M13 peptide, enabling hinge movements of its two domains (Figure 1C, D illustrates this movement). Several hinge points have been examined in nearby (nonadjacent) residues, still reproducing similar configurations which are in general agreement with the native complexed structure. The average RMSD of interface atoms of the correct solution is around 5 Å, and average matching time is 0.25 minutes (see columns 8 and 9 in Table 4).

Analysis of the results have shown that there is a correlation between the contact percentage which is used for ranking the acceptable binding modes and the voting (matching) score which is based on the number of votes received by the computed candidate hinge locations. The best solution (low RMSD) usually scores high as a binding mode having a good fit with relatively high contact percentage. It can be found between the first top scoring solutions and the 9% ones. Its voted for hinge location may be identified at the end of the matching stage, since the transformed hinge location of the best solution usually receives a large number of votes. The docked configurations we have obtained may provide initial guesses for subsequent detailed conformational space sampling (around the hinge-bent configurations) or for molecular dynamics simulations. Energy minimization steps may follow to refine and optimize the configurations.

4. DISCUSSION

In this work, we present a general approach to docking a ligand onto a receptor, allowing hinge-bending motions of relatively rigid parts in either the ligand or receptor molecules. Our main algorithmic tool is the generalization and the extension of the generalized Hough transform technique, which was originally developed for partially occluded articulated object recognition in Computer Vision and Robotics. A major problem in Computer Vision is object recognition in cluttered scenes, where the objects may be partially occluded. The recognition problem becomes especially challenging when the objects are so-called, articulated objects—namely, objects consisting of rigid parts connected by either rotary or sliding joints. Efficient techniques to handle these types of problems have been suggested (Wolfson, 1991). From a geometric standpoint, there exists an analogy between the problems of object matching and assembly in Computer Vision and Robotics, and those of molecular structural comparison and docking in Molecular Biology. In both cases, one seeks to discover subparts of the objects (molecules) under study, which have a similar geometric structure (molecular surfaces patches). In general, in both cases, one does not have an *a priori* knowledge of the subpart (binding site) that will exhibit the match. The analogy between these types of problems brought about this interdisciplinary research endeavor (for previous related research, see Nussinov and Wolfson, 1991).

Considering the implications of our algorithm to the recognition problems in Computer Vision, our method has the following attributes:

1. *A technique for addressing 3-D to 3-D matching problems.* We have extended the 2-D generalized Hough transform based technique to 3-D, thus potentially enabling the handling of occluded articulated object recognition problems originating from diverse research fields.
2. *A full 3-D joint rotation of the object during matching.* This general geometrical model enables all rotary trajectories possible by point rotation in space. It is not restricted to one or two degrees of freedom rotation around an axis (axes). The latter restriction is optionally applied during the subsequent verification stage.
3. *Efficient handling of a large number of interest features.* Usually, algorithms in Computer Vision handle a small number of interest features specifying the objects (typically 10–20 interest features). Our matching algorithm efficiently handles objects consisting of hundreds of interest features, as well as objects composed of tens of interest points (we have matched objects consisting of 33–804 feature points). Typical matching




FIG. 6. Examples of the MTX/DHFR geometrically acceptable binding conformations. The ligand molecule is drawn in its van der Waals solid sphere representation. The DHFR molecule is represented as a white ribbon. The location and orientation of the DHFR is the same in all figures. (a) The original crystal bound complex. (b) A predicted binding mode of the MTX ligand. (c) An additional predicted binding mode of the MTX ligand.

run-times are a few seconds to around half a minute, depending on the number of interest features of the matched objects (run-time measurements have been conducted on a SGI R10000 machine).

4. *Handling of noisy sampled input within cluttered scenes.* Due to the complex nature of the molecular structure determination process, the input data to our algorithm can have a range of resolutions and accuracies. Our method yields reasonable results, manifesting robustness and tolerance in tackling this difficulty. Also, our method handles occluded objects in highly cluttered scenes.
5. *Complexity of the matching is $O(n^3 \times R)$* , where n is the number of the target interest features. The power of 3 represents the object's shape signature (surface patches) complexity. The shape signature is defined as a triplet of points, allowing a unique definition of a 3-D rotation and translation, which superimposes a model structure to the target structure. The model objects are indexed into a look-up table (hash-table), allowing the target scene to be simultaneously matched with all previously stored model(s). R is the average number of records within the associative look-up table entries, i.e., the average "access time" to the table. Constant value of R may result in a practical complexity of order n^3 .
6. *Incorporation of the information from all parts of the objects, in the recognition process, regardless of the part's size.* We exploit the fact that the different parts belong to the same object and share common joints. There is no order in the recognition of the parts, and no additional degrees of freedom are incorporated in the matching process. Thus, the parts are not restricted to being highly informative, and the global consistency checks are conducted within the recognition process itself.

The capabilities of our algorithm, manifested in the Molecular Biology applications presented here, have several attractive features.

1. *The hinge is introduced either into a large receptor molecule allowing domain motions, or into the ligand.* To date, other methods allow hinges to be positioned in small ligands, and only partial (e.g., side chain) flexibility is permitted in the receptors.
2. *Docking of diverse sized molecules.* Our general method allows hinge induced motions to exist in either variable size receptors, or in diverse size ligands. The sizes of the molecules range from tens of atoms to thousands of atoms (we have docked molecules of 23–2,862 atoms).
3. *Good quality predictive results.* The correct conformations are obtained with small RMS deviations from the experimental (crystal/NMR determined) configurations. Additional predictive, high-scoring (good-fitting) binding modes, are generated as well.
4. *A substantial conformational change between the bound and unbound structures is handled,* yielding compatible and acceptable docked configurations.
5. *Short execution times for matching the molecular surfaces.* Seconds to around half a minute, mostly depending on the target molecule size.
6. *Decrease in running times* for the incorporation of additional hinges.
7. *Entire molecular surfaces are considered,* assuming no knowledge of the binding site.
8. *Unbound and bound NMR and crystal molecular structures are successfully docked.* In general, docking NMR structures presents a higher degree of difficulty, as they are determined from an ensemble of relatively flexible conformations present in solution.

It is the combination of the above that enables the application of this tool to molecular motions which are predominantly hinge. However, there are some limitations. If the molecular parts are relatively small, e.g., a part comprising a single bond, they may not be efficiently matched by the method, as small subparts are described by a small number of interest features. A plausible remedy for this type of cases, is to explore the conformational space. In addition, when dealing with molecular databases, the data structures employed here require a significant memory space. Thus, when extending the method to handle database screening, an alternative storage areas should be considered, e.g., external, direct access storage.

Our ability to predict good fitting binding modes of the docked conformations, provides a useful research tool in biomolecular structural studies and can aid in rational drug design. The association between the receptor and the ligand is an essential step in many biological processes, such as chemotherapy, regulatory mechanisms and toxicity. Knowing the 3-D structure of the associating molecules, our ability to predict their molecular interactions can potentially lead to the discovery and to the design of a chemical or a compound specifically fitting a target molecule. Structure-based drug design is an increasingly promising approach for rational drug development. Many therapeutic agents have been recently spawned by this approach, for treating diseases such as cancer and AIDS. In comparison with the traditional approaches, automation can provide the speed-up




FIG. 7. The *unbound* case where the hinge is in the unbound maltose-binding protein (MBP) receptor and the maltose ligand is rigid. We position the hinge at the C α atom of the Glu¹¹¹, pointed at by the arrow. The maltose ligand is represented as van der Waals solid spheres. The MBP receptor is plotted as a ribbon. (a) The unbound MBP receptor and the maltose ligand extracted from two different PDB files. The MBP is in its "open conformation." (b) The unbound MBP receptor is docked with the maltose ligand yielding a "close conformation" of the receptor, which engulfs the ligand.

required for large-scale screening efforts of drug databases. For example, the Fine Chemical Directory (FCD) database (Gunner *et al.*, 1991) and the Chemical Structural Database (CSD) (Allen *et al.*, 1979). In the classical methods, several thousands of compounds may have to be synthesized and tested, before a suitable drug is found (Branden and Tooze, 1991). The structure-based methodology can yield drugs that are "tailored" for their target, more quickly and less expensively than other ways. Owing to these advantages, the compounds can potentially be more specific, more potent, and less toxic than those discovered using other methodologies (Bugg *et al.*, 1993; Navia and Murcko, 1992; Martin, 1991; Kuntz, 1992). An efficient method for molecular docking, as presented here, geared toward screening molecular databases, can contribute to achieving these goals. Many compounds of pharmacological interest can be simultaneously screened for docking to a target molecule. The flexibility we allow, in either the ligand or receptor molecules, can yield additional therapeutic agents that otherwise are most likely to be ignored and missed when using the more restrictive rigid body docking techniques.

We have obtained the correct binding modes which are consistent with experimental results by successfully reproducing known docked configurations of the associating molecules. The alternate docked configurations that we predict, i.e., the additional good-fitting binding modes between the two molecules, provide predictions for plausible molecular interactions and conformations. This capability is important to biomolecular research and to structure-based drug design, in the case where the 3-D molecular structures have been elucidated. Also, the predictive binding modes are important for investigations of molecular interactions. The geometrically acceptable docked solutions can be input to routines examining the chemical interactions between the receptor-ligand atoms at the interface. Biochemical feasibility may be checked in terms of, for example, electrostatic forces, hydrogen-bonding, and hydrophobicity, thus discriminating between the biologically favorable and unfavorable solutions.

Following are some additional features that may be incorporated into our algorithm for improving its efficiency and execution times. The algorithm presented here may be extended to account for chemical filtering during the matching process. This may be done, for example, by labeling the "interest points" according to chemical properties, and accepting only geometrically and chemically consistent matches (Rarey *et al.*, 1996). Not only will this procedure obviate the need for carrying out chemical verification of all geometrically acceptable docked conformations, but it would also filter out many biologically unacceptable solutions, speeding up the verification stage. Furthermore, since rotational bond movement is more restrictive than a 3-D hinge rotation, this restriction may be incorporated in the algorithm (matching stage) and improve its performance for the cases where a full 3-D point rotation is not required (see, for example, Rigoutsos and Califano, 1996). Parallelization of the algorithm may also speed up the calculation.

The generality of our approach allows its application toward additional structural/pattern matching problems in molecular biology, such as protein structure comparisons. These can be exemplified by a database search for hinge-bent motifs, which can aid in investigations of protein folding. Moreover, our algorithm can be applied to recognition problems in Computer Vision, such as problems arising from diverse fields of medical imaging, CAD/CAM, robotics assembly, mobility, and manipulation.

ACKNOWLEDGMENTS

We thank Drs. Shuo L. Lin, Robert Jernigan, and Jacob Maizel, for helpful discussions, encouragement and interest. The research of R. Nussinov has been sponsored by the National Cancer Institute, DHHS, under contract no. 1-CO-74102 with SAIC. The contents of this publication do not necessarily reflect the views or policies of the DHHS, nor does mention of trade names, commercial products, or organization imply of endorsement by the U.S. Government. The research of H. J. Wolfson and R. Nussinov in Israel has been supported in part by grant no. 95-0028 from the U.S.-Israel Binational Science Foundation (BSF), Jerusalem, Israel, by the Recanati fund, and by a basic research fund from Tel Aviv University. H.J.W. acknowledges the support of the Hermann Minkowski-Minerva Center for Geometry at Tel Aviv University. This work forms part of the doctoral dissertation of B. Sandak, Tel Aviv University, and is supported in part by the Carl F. Gauss Minerva Center for Scientific Computation at the Weizmann Institute of Science.

REFERENCES

- Allen, F., Bellard, S., Brice, M., *et al.* 1979. The Cambridge Crystallographic Data Center: computer-based search, retrieval, analyses and display information. *Acta Crystallogr. B* 35, 2331-2339.

- Ballard, D. 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13(2), 111–122.
- Beinglass, A., and Wolfson, H. 1991. Articulated object recognition, or, how to generalize the generalized Hough transform, 461–466. In *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*. Maui, Hawaii.
- Bernstein, F., Koetzle, T., Williams, G., et al. 1977. The protein data bank: a computer-based archival file for macromolecular structures. *J. Mol. Biol.* 112, 535–542.
- Bolin, J., Filman, D., Matthews, D., Hamlin, R., and Kraut, J. 1982. Crystal structures of *Escherichia coli* and *Lactobacillus casei* dihydrofolate reductase refined at 1.7 Å resolution. *J. Biol. Chem.* 257, 13650–13662.
- Branden, C., and Tooze, J. 1991. *Introduction to Protein Structure*. Garland Publishing, New York.
- Bugg, C., Carson, W., and Montgomery, J. 1993. Drugs by design. *Sci. Am.* Dec, 92–98.
- Clark, K., and Ajay 1995. Flexible ligand docking without parameter adjustment across four ligand receptor complexes. *J. Comp. Chem.* 16, 1210–1226.
- Connolly, M. 1983a. Solvent-accessible surfaces of proteins and nucleic acids. *Science* 221, 709–713.
- Connolly, M. 1983b. Analytical molecular surface calculation. *J. Appl. Cryst.* 16, 548–558.
- DesJarlais, R., Kuntz, I., Furth, P., et al. 1990. Structured-based design of nonpeptide inhibitors specific for the human immunodeficiency virus-1 protease. *Proc. Natl. Acad. Sci. U.S.A.* 87, 6644–6648.
- DesJarlais, R., Sheridan, R., Dixon, J., Kuntz, I., and Venkataraghavan, R. 1986. Docking flexible ligands to macromolecular receptors by molecular shape. *J. Med. Chem.* 29, 2149–2153.
- Fischer, D., Lin, S.L., Wolfson, H.J., and Nussinov, R. 1995. A geometry-based suite of molecular docking processes. *J. Mol. Biol.* 248, 459–477.
- Ghose, A., and Crippen, G. 1985. Geometrically feasible binding modes of a flexible ligand molecule at the receptor site. *J. Comp. Chem.* 6, 350–359.
- Goodsell, D., and Olson, A., 1990. Automated docking of substrates to proteins by simulated annealing. *Proteins Struct. Funct. Genet.* 8, 195.
- Gunner, O., Hughes, D., and Dumont, L. 1991. *Chem. Inf. Comput. Sci.* 31, 408.
- Halperin, D., and Overmars, M. 1994. Spheres, molecules, and hidden surface removal, 113–122. In *Proceedings of the 10th ACM Symposium on Computational Geometry*. Stony Brook.
- Jiang, F., and Kim, S. 1991. Soft docking: matching of molecular surface cubes. *J. Mol. Biol.* 219, 79–102.
- Jones, G., Willet, P., and Glen, R. 1995. Molecular recognition of receptor sites using a genetic algorithm with a description of desolvation. *J. Mol. Biol.* 245, 43–53.
- Jones, G., Willett, P., Glen, R., Leach, A., and Taylor, R. 1997. Development and validation of a genetic algorithm for flexible docking. *J. Mol. Biol.* 267, 727–748.
- Katchalski-Katzir, E., Shariv, I., Eisenstein, M., Friesem, A., Aflalo, C., and Vakser, I. 1992. Molecular surface recognition: determination of geometric fit between proteins and their ligands by correlation techniques. *Proc. Natl. Acad. Sci. U.S.A.* 89, 2195–2199.
- Kuntz, I. 1992. Structure-based strategies for drug design and discovery. *Science* 257, 1078–1082.
- Kuntz, I., Blaney, J., Oatley, S., Langridge, R., and Ferrin, T. 1982. A geometric approach to macromolecule-ligand interactions. *J. Mol. Biol.* 161, 269–288.
- Lamdan, Y., Schwartz, J., and Wolfson, H. 1990. Affine invariant model-based object recognition. *IEEE Trans. Robot. Automat.* 6(5), 578–589.
- Leach, A. 1994. Ligand docking to proteins with discrete side-chain flexibility. *J. Mol. Biol.* 235, 345–356.
- Leach, A., and Kuntz, I. 1992. Conformational analysis of flexible ligands in macromolecular receptor sites. *J. Comp. Chem.* 13(6), 730–748.
- Lin, S.L., and Nussinov, R. 1996. Molecular recognition via the face center representation of molecular surface. *J. Mol. Graphics* 14, 78–97.
- Lin, S.L., Nussinov, R., Fischer, D., and Wolfson, H. J. 1994. Molecular surface representation by sparse critical points. *Proteins Struct. Funct. Genet.* 18, 94–101.
- Martin, Y. 1991. Computer-assisted rational drug design. *Methods Enzymol.* 203, 587–613.
- Mizutani, M., Tomioka, N., and Itai, A. 1994. Rational automatic search method for stable docking models of protein and ligand. *J. Mol. Biol.* 243, 310–326.
- Navia, M., and Murcko, M. 1992. Use of structural information in drug design. *Curr. Opin. Struct. Biol.* 2, 202–210.
- Nussinov, R., and Wolfson, H.J. 1991. Efficient detection of three-dimensional motifs in biological macromolecules by computer vision techniques. *Proc. Natl. Acad. Sci. U.S.A.* 88, 10495–10499.
- Oh, B., Pandit, J., Kang, C., et al. 1993. Three-dimensional structures of the periplasmic lysine/arginine/ornithine-binding proteins with and without a ligand. *J. Biol. Chem.* 268, 11348–11355.
- Rarey, M., Kramer, B., Lengauer, T., and Klebe, G. 1996. A fast flexible docking method using incremental construction algorithm. *J. Mol. Biol.* 261, 470–489.
- Richards, F. 1977. Areas, volumes, packing and protein structure. *Annu. Rev. Biophys. Bioeng.* 6, 151–176.
- Rigoutsos, I., Platt, D., and Califano, A. 1996. Flexible 3D-substructure matching and novel conformer derivation in very large databases of 3D-molecular information. Research Report RC 20497 (90854), IBM Research Division.
- Sandak, B. 1997. A Method for Biomolecular Structural Recognition Allowing Conformational Flexibility [Dissertation]. Computer Science Department, Tel Aviv University.

- Sandak, B., Nussinov, R., and Wolfson, H.J. 1995. An automated computer-vision and robotics-based technique for 3-D flexible biomolecular docking and matching. *Comput. Appl. Biosci.* 11, 87–99.
- Sandak, B., Nussinov, R., and Wolfson, H.J. 1996a. Docking of conformationally flexible proteins. *Lect. Notes Comput. Sci.* 1075, 271–287.
- Sandak, B., Wolfson, H.J., and Nussinov, R. 1996b. Hinge-bending at molecular interfaces: automated docking of a dihydroxyethylene-containing inhibitor of the HIV-1 protease. *J. Biomol. Struct. Dyn.* 1, 233–252.
- Sharff, A., Rodseth, L., Spurlino, J., and Quioco, F. 1992. Crystallographic evidence of a large ligand-induced hinge-twist motion between the two domains of the maltodextrin binding protein involved in active transport and chemotaxis. *Biochemistry* 31, 10657–10663.
- Smellie, A., Crippen, G., and Richards, W. 1991. Fast drug-receptor mapping by site-directed distances: a novel method of predicting new pharmacological leads. *J. Chem. Inform. Comput. Sci.* 31, 386–392.
- Spurlino, J., Lu, G., and Quioco, F. 1991. The 2.3Å resolution structure of the maltose or maltodextrin binding protein, a primary receptor of bacterial transport and chemotaxis. *J. Biol. Chem.* 266, 5202–5219.
- Theriault, Y., Logan, T., Meadows, R., *et al.* 1993. Solution structure of the cyclosporin A/cyclophilin complex by NMR. *Nature* 361, 88–94.
- Wang, H. 1991. Grid-search molecular accessible algorithm for solving the protein docking problem. *J. Comp. Chem.* 12, 746–750.
- Welch, W., Ruppert, J., and Jain, A. 1996. Hammerhead: fast, fully automated docking of flexible ligands to protein binding sites. *Chem. Biol.* 3, 449–462.
- Wolfson, H.J. 1991. Generalizing the generalized Hough transform. *Pattern Recognition Lett.* 12(9), 565–573.

Address reprint requests to:

Bilha Sandak
Department of Applied Mathematics and Computer Science
Weizmann Institute of Science
Rehovot 76100
Israel

billie@wisdom.weizmann.ac.il

Received for publication August 6, 1997; accepted as revised May 3, 1998.