CABIOS

Vol. 11 no. 1 1995
Pages 87–99

An automated computer vision and robotics-based technique for 3-D flexible biomolecular docking and matching

Bilha Sandak[1], Ruth Nussinov[2,3]*, and Haim J.Wolfson[1]

Abstract

The generation of binding modes between two molecules, also known as molecular docking, is a key problem in rational drug design and biomolecular recognition. Docking a ligand, e.g., a drug molecule or a protein molecule, to a protein receptor, involves recognition of molecular surfaces as molecules interact at their surface. Recent studies report that the activity of many molecules induces conformational transitions by 'hinge-bending', which involves movements of relatively rigid parts with respect to each other. In ligand-receptor binding, relative rotational movements of molecular substructures about their common hinges have been observed. For automatically predicting flexible molecular interactions, we adapt a new technique developed in Computer Vision and Robotics for the efficient recognition of partially occluded articulated objects. These type of objects consist of rigid parts which are connected by rotary joints (hinges). Our approach is based on an extension and generalization of the Geometric Hashing and Generalized Hough Transform paradigm for rigid object recognition. Unlike other techniques which match each part individually, our approach exploits forcefully and efficiently enough the fact that the different rigid parts do belong to the same flexible molecule. We show experimental results obtained by an implementation of the algorithm for rigid and flexible docking. While the 'correct', crystal-bound complex is obtained with a small RMSD, additional, predictive 'high scoring' binding modes are generated as well. The diverse applications and implications of this general, powerful tool are discussed.

Introduction

The problem of generating feasible binding modes between two molecules, is referred to as the *molecular docking* problem. The ability to *automatically* predict molecular interactions, is important in rational drug design and discovery, as well as a research tool in biomolecular structural recognition. Geometrical fitness is a necessary condition for successfully docking molecules. The docking problem is thus usually approached by first determining the geometrically acceptable solutions and then checking the proposed bound complexes for their chemical/physical/biological feasibility. As molecules interact at their surface, molecular docking involves *recognition* of molecular surfaces. By considering molecules as 3-dimensional (3-D) structures, specified by their molecular surface representation, techniques originated in Computer Vision and Robotics can be applied to discover docked solutions. The recognition process we face in the docking problem is reminiscent of the automatic part assembly problem in robotics and, if one of the molecules is taken as a complement, is equivalent to the 3-D partially occluded object recognition task in Computer Vision. This task is usually formulated as follows : Given a data base of previously known objects and a newly observed scene with numerous cluttered objects, recover all the occurrences of the database objects in the scene, even if they are partially occluded. The analogy is quite obvious. The database of objects becomes a database of ligands (e.g., drug molecules or protein molecules) and the newly observed scene is the 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 about the location of the binding site.

Rigid docking methods consider the molecules as 3-D rigid structures. They are aimed at rigid docking of a ligand to a rigid protein receptor. The former may be regarded as the 'key' and the latter as the 'lock'. Rigid-body approaches, even if successful (Jiang and Kim, 1991; Wang, 1991), and efficient (Fischer *et al.*, 1993) address a particular case. Yet, these methods do not take into account the conformational changes that molecules may undergo. Here we present a general docking approach, allowing *hinge-bending* of relatively rigid parts. These movements are characterized by the rotational movements of molecular substructures about their hinges. Hinge movements of molecular domains have been observed in immunoglobulins (Bennet and Haber, 1984), T4 lysozyme (Faber and Matthews, 1990; Dixon *et al.*, 1992), lactate dehydrogenase (Gerstein and Chothia, 1991), aspartic

[1]*Computer Science Department, School of Mathematical Sciences, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 69978, Israel,* [2]*Laboratory of Mathematical Biology, PRI/Dynacorp, NCI-FCRF, Bldg 469, rm 151, Frederick, MD 21712, USA,* [3]*Sackler Institute of Molecular Medicine, Faculty of Medicine, Tel Aviv University, Tel Aviv 69978, Israel*

**Corresponding author (at USA)*

proteinases (Sali *et al.*, 1992), and other molecules. For in depth review of hinge based domain movements see (Gerstein *et al.*, 1994). Hinge bending has also been observed in ligand-receptor binding. The binding of cyclosporin-A to cyclophilin is an example of a flexible ligand that undergoes geometrical distortion to achieve a suitable binding conformation (Weber *et al.*, 1991). Conformational changes in antigen-antibody binding have been reported in (Rini *et al.*, 1992; Stanfield *et al.*, 1990). Flap domain motions in the binding of HIV-1 protease and a peptide inhibitor have been observed by (Miller *et al.*, 1989). The computational difficulty grows considerably when taking into account the additional degrees of freedom inherent in the *flexible molecular docking* problem.

We consider the molecules as flexible 3-D structures and the flexibility we allow is rotational movements of the molecular substructures about hinges. The close analogy between rigid docking and robotic assembly extends also to the flexible case, which is usually referred to in Computer Vision and Robotics as *articulated object recognition (matching)*. Articulated objects are objects consisting of rigid parts which are connected either by rotary or sliding (prismatic) joints. The analog of a hinge is a rotary joint. Up to date very little work was done on articulated object matching in Computer Vision (Brooks, 1981; Grimson, 1991; Lowe, 1991; Shakunaga, 1991; Wolfson, 1991). Even less work was done on matching flexible structures in molecular biology (DesJarlais *et al.*, 1986; Leach and Kuntz, 1991). Other methods which account for conformational flexibility are the distance geometry method of (Ghose and Crippen, 1985) and the simulated annealing approach of (Goodsell and Olson, 1990). These latter methods are characterized by randomly generating conformations and verifying their feasibility. A typical approach to tackle the articulated object recognition problem in Computer Vision (Grimson, 1991) and in Molecular Biology (DesJarlais *et al.*, 1986; Leach and Kuntz, 1992) is to represent an object as a composition of its rigid parts, try to match each part individually, and then check some global consistency among the candidate solutions for the separate parts. This type of approach is conceptually identical to the rigid case, and, if the individual parts do not have a large matching area, there is insufficient information to proceed. One can say that this approach does not exploit forcefully enough the fact that the different parts do belong to the same object.

We suggest a flexible docking approach which incorporates both the more simple rigid subpart recognition techniques and the global consistency checks as an integral part of the recognition process. Our approach is inspired by our previous work in Computer Vision (Wolfson, 1991), which extended the rigid body matching techniques based on Geometric Hashing (Lamdan *et al.*, 1990) and on the Generalized Hough Transform (Ballard, 1981). A preliminary implementation of this approach for Computer Vision applications of industrial tool recognition in photographs has been reported in (Beinglass and Wolfson, 1991). Here we report our implementation of this approach for the 3-D molecular docking case.

The algorithm

We present the algorithm in the case of one hinge in the ligand. Although hinge-movements can be introduced either into the receptor or into the ligand, or into both, we position the hinge in the ligand. The considerably larger receptor is assumed to be rigid. We base our choice on studies that show that the smaller ligands are considerably more flexible than the larger receptors. We describe our algorithm for the case of docking ligands which consist of two rigid substructures connected by a rotary hinge (joint), as illustrated in Figure 1. A similar technique applies for the case of *multiple hinges* as well. To simplify the exposition we focus on the single hinge case. At the end of this section we briefly sketch the modifications needed for the multiple hinge case. We allow a full 3-D rotation at the hinge. In practice, rotational bonds have only a one degree of freedom rotation around a known axis (the bond). The algorithm that we present can be easily modified for this simpler case resulting in increased time efficiency and fewer 'false positives'. Nevertheless, we preferred to conduct our initial investigations with a more general rotational model, namely, the one with 3 degrees of freedom. As a by-product, our geometric model allows a good approximation to the case of two consecutive rotational bonds.

We base our approach on a 'voting scheme' for finding the most suitable ligands (out of a library of ligands), and the transformations for their docking with a target receptor. First, both the ligand and the receptor molecules are reduced to a certain number of 'interest points'. This reduction is described below and in (Fischer *et al.*, 1993). These point sets are then handled by our geometric matching algorithm. Thus, the matching is done between these 3-D sets of 'interest points'. However, while in (Fischer *et al.*, 1993) these sets have been assumed to be rigid, here we allow a rotary hinge in the ligand.

As ligands may undergo translations and rotations of the parts in order to dock to a receptor, the ligand description is stored in a database invariant to this type of transformations. This procedure is carried out in the *preprocessing* step of the algorithm. The location of the hinge for each ligand stored is assumed to be known in this step. A receptor structure is presented to the system in the

recognition step of the algorithm. If a ligand, previously stored in the database, has a matching surface patch (i.e., an 'interest point' configuration) which has similar geometric attributes to a receptor surface patch, a vote is cast for this ligand together with the computed location of its hinge. We define an ordered non-collinear triplet of points as an 'interest point' configuration. Corresponding ligand and receptor ordered triplets uniquely define a 3-D transformation (rotation and translation). The computed hinge location is derived from the transformation between the corresponding receptor and ligand patches. Exploiting the fact that both parts of a ligand incorporate the same hinge, although being at different orientations, both parts can potentially contribute matching 'votes' to the same hinge location (see Figure 1). A description of our two basic-step algorithm is described below. The scheme of the algorithm is presented in Figure 2.

Preprocessing–where the description of the ligands is stored in the database, that is, the ligand information is indexed into a look-up (hash) table. Since this step is independent of the receptor structure, it is executed off line.

Recognition–where a new receptor structure is introduced. The structures of those ligands having relatively large matching section(s) with it, are recovered from the look-up table. Transformations of the ligands-parts yielding matches with the receptor are obtained (*matching stage*). The transformations are filtered for obtaining acceptable solutions. This is carried out by rejecting candidate transformations which result in collision between the ligand's two parts or between the receptor and the ligand (*verification stage*).

Preprocessing

The preprocessing procedure is given below. See also Figure 3.

Procedure *Preprocessing*,
  Do for all ligands
    Extract the ligand molecule 'interest points' (atoms),
    Pick the rotary hinge as the ligand's 'reference center',
    Define and center at the hinge a 3-D orthonormal coordinate frame,
      referred to as the 'hinge-based coordinate frame',
    Do for all ordered, non-collinear triplets of 'interest points'·
      Define three unique orthonormal coordinate frames based on a triplet,
        referred to as 'triplet-based coordinate frames',
      Compute an entry address in the look-up table which is the
        transformation-invariant three inter-point distances,
      Compute and store the transformations between the 'triplet-based coordinate frames'
        and the 'hinge-based coordinate frame' in a look-up table entry record,
      Store the ligand identification at the same record of the entry,
    End,
  End,
End.

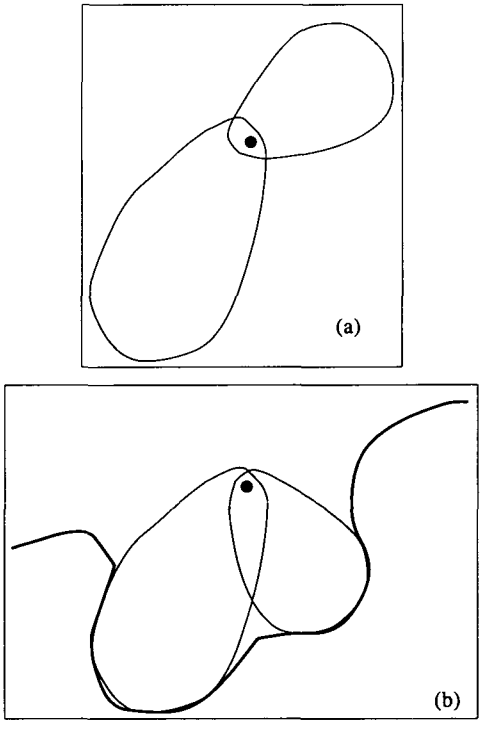

Fig. 1. A schematic illustration of hinge-bending. (**a**) A flexible ligand consisting of two rigid parts connected by a hinge. (**b**) The flexible ligand displayed in (a), is docked to a receptor The two rigid parts having matching surface areas with the target structure, incorporate the hinge at the same location, although the two parts are at different orientations.

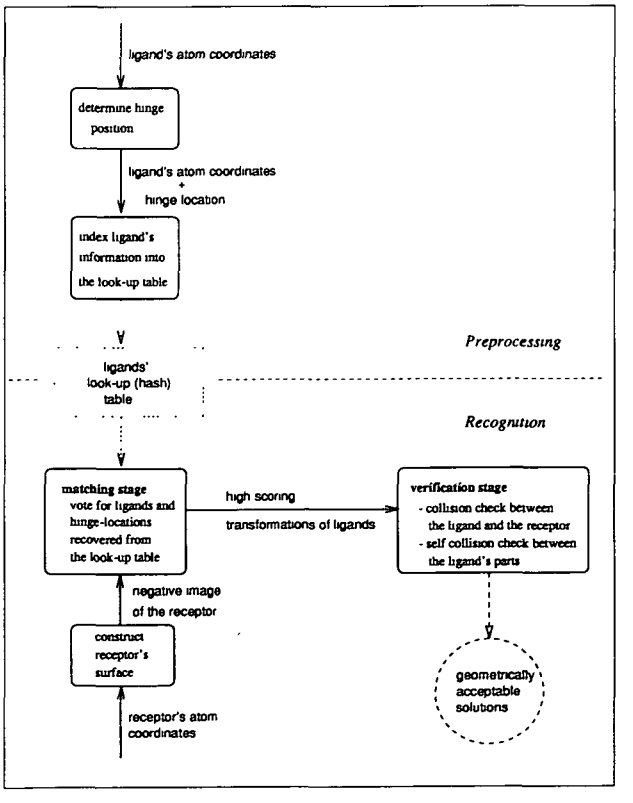

Fig. 2. A general scheme of the algorithm.

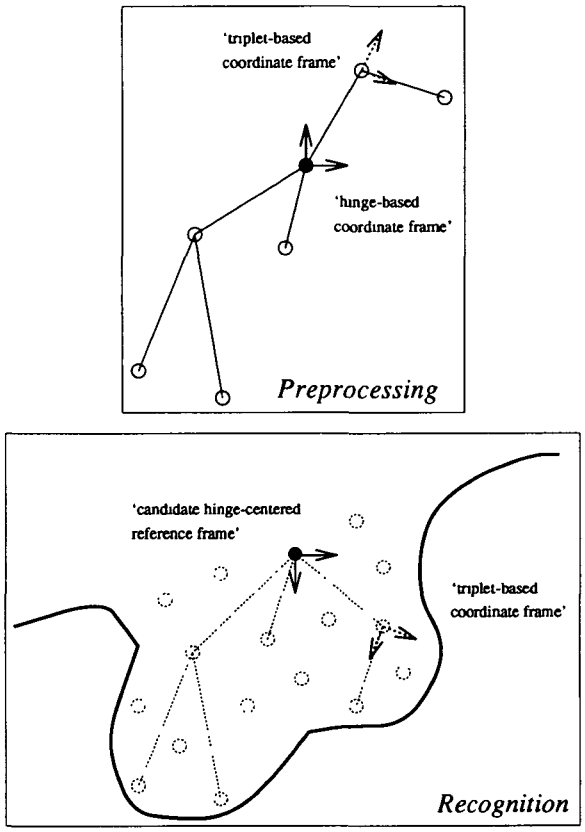

Fig. 3. A 2-D illustration of the two-basic steps of the algorithm First step—*preprocessing*: A flexible ligand consisting of two rigid sub-parts connected by a hinge. The small circles represent the ligands atoms; the lines represent the covalent bonds, which are disregarded by our algorithm. The hinge (dark circle) is picked as the ligand's 'reference center'. An orthonormal coordinate frame is defined and centered at this hinge—referred to as the 'hinge-based coordinate frame'. For each (non-collinear) triplet of atoms, we define an orthonormal 'triplet-based coordinate frame'. We compute and store the transformation between the 'triplet-based coordinate frame' and the 'hinge-based coordinate frame' in the look-up (hash) table If there are several ligands which will be docked to the receptor, the ligand's identification (name) is stored in the table as well Second step—*recognition*· The flexible ligand presented above is docked onto the receptor surface. The small circles represent the sphere centers of the negative image of the receptor wherein the ligands, previously stored in the look-up table, are to be recognized. The two rigid parts having matching sections with the receptor structure, incorporate a reference frame, a 'candidate hinge-centered reference frame', at the same location (dark circle), although at two different orientations. The 'candidate hinge-centered reference frames' are calculated by applying the pre-recorded transformations of the 'triplet-based coordinate frames' of the ligand, stored in the look-up table, to all the 'triplet-based coordinate frames' of the receptor.

The complexity of the preprocessing step is of order $N \times m^3$, where m is the number of atoms in the ligand, and N, the number of ligands in the database. The addresses of the entries of the table are discretized into bins. In our implementation, the bins are of 0.5 Å. Clearly, each table entry may contain more than one description record of a ligand's surface patch.

Recognition

The molecular structure of the receptor, is represented by its negative (complementary) image. The complementary shape of the receptor is generated by constructing the molecular surface (Connolly, 1983a; Connolly, 1983b) and detection of the invaginations on the receptor surface (Kuntz *et al.*, 1982). The negative image of the receptor is represented as clusters of intersecting spheres. The largest of the clusters is considered to be the binding site of the ligand. For small-molecule ligands, the prediction of the binding site on the receptor is quite accurate. This clustered sphere representation, i.e., the coordinates of the sphere centers, are considered the receptor-scene wherein section(s) of a matching ligand 'interest point' configuration is (are) to be recognized. The recognition procedure and its stages are described in detail below. See also Figure 3 for a graphical representation of the matching stage.

Procedure *Recognition*,
  Extract the receptor molecule 'interest points' (sphere centers),
  Call the procedure *matching_stage*;
  Call the procedure *verification_stage*,
End,

Procedure *matching_stage*,
  Do for all non-collinear triplets of 'interest points'
    Define three unique orthonormal coordinate frames based on a triplet,
      referred to as 'triplet-based coordinate frames',
    Compute an entry address in the look-up table which is the
      transformation-invariant three inter-point distances;
    Do for each record in the resulting look-up table entry
      Apply the pre-recorded transformations to the 'triplet-based coordinate frames'
        and compute the 'ligand hinge locations' and 'reference frames',
      Cast a vote for the identity of the ligand together with the
        location and orientation of the 'candidate hinge-centered reference frames',
    End,
  End;
End,

After the matching stage, we are left with many relatively high scoring candidate solutions. We filter the transformations, having good matches with the receptor, to obtain acceptable solutions. This procedure is carried out in the verification stage, presented below. A potential match implies existence of complementarity between receptor-ligand surface patches. However, other regions of the two molecules may collide. In the *collision check* we reject the transformations which result in collision between the ligand and the receptor. The transformations 'surviving' the collision check are then passed to the *self collision check*. In this procedure, we discard the candidate solutions which result in collision between the two ligand parts.

Procedure *verification_stage*;
  Call the procedure *collision_check*,
  Call the procedure *self_collision_check*;
End;

Since we are interested only in high scoring transformations, we pick hinge locations receiving a large number of

votes. 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. The high scoring transformations are filtered to obtain transformations which result in docking the ligand onto the receptor with no penetration between the ligand and the receptor, and the ligand's parts. 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 two ligand parts. The van der Waals radii of the current receptor atom and ligand atom checked are denoted by $R_{receptor\_atom}$ and R_{ligand_atom}, respectively. To speed up the collision check we reduce the 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). 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 Walls radii plus a proximity threshold (*contact distance*).

```
Procedure collision_check,
  'Divide' the receptor molecule into eight segments,
  Search the accumulator of votes for high scoring pairs (ligand, hinge location),
      determined according to the voting threshold,
  /* do the following for all high scoring ligands */
  Do for all high scoring hinge locations (for each part of the ligand individually)
    Compute the transformation between the 'hinge-based coordinate frame'
    and the 'candidate hinge-centered reference frame'
    Do for each atom in the respective part of the ligand
      Apply the computed transformation to the atom,
      Determine the receptor segment the transformed atom is in,
      Do for each receptor atom in the segment·
        Compute the distance d between the ligand atom and the receptor atom,
        If |d| < (R_ligand_atom + R_receptor_atom) − collision_distance then
          Ligand part and receptor collide → Discard the transformation,
          Goto next computation of transformation,
        else
          If |d| < (R_ligand_atom + R_receptor_atom) + collision_distance then
            Mark the atom as being in contact with the receptor,
      End,
    End;
    Compute the contact percentage of the ligand's part according to the marked atoms,
  End,
End;
```

As mentioned above, the self collision check employs the same criterion for rejecting self penetration causing transformation, as being done by the collision check. The van der Waals radii of the currently checked atoms belonging to the different parts, are denoted by R_{part1_atom} and R_{part2_atom}, respectively.

```
Procedure self_collision_check,
  /* do the following for all high scoring ligands */
  Do for every hinge location (which is actually the 3-D translation of the ligand),
    Do for every rotation of the first part;
      Do for every rotation of the second part,
        Do for each atom in first part,
          Apply the transformation (translation + rotation of the part)
          to the atom,
          Do for each atom in second part;
            Apply the transformation (translation + rotation of the part)
            to the atom,
            Compute the distance d between the atoms,
            If |d| < (R_part1_atom + R_part2_atom) − collision_distance then
              Ligand's parts collide → Discard the transformations;
              Goto next transformation;
          End;
        End;
        Compute the contact percentage of the ligand by summing
        the contact percentage of its parts,
      End;
    End;
  End,
End,
```

In this algorithm we have exploited the fact that both parts incorporate the same hinge by locating the reference frame of the ligand at the hinge. In such a way both parts contribute votes to a reference frame at the same location, although at different orientations. By picking up votes from both of its parts, a ligand, which might have a small matching surface area with the receptor structure in each of its parts, can still score high, although each of its individual parts could receive an insignificant score.

The complexity of our algorithm is of order $n^3 \times A$, where n is the number of receptor surface points considered in the matching, and A is the access time to a record in an entry in the look-up table. 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 ignores extremely large bins, this access time will be constant and independent of n, thus achieving practical complexity of the order of n^3.

The flexible docking algorithm can handle the rigid docking as a particular case. In the rigid-body docking, the ligand's reference frame is located arbitrarily.

Although here, our algorithm is applied to molecules with a single hinge, it can be extended to multiple hinges by the following modification. Instead of having one 'hinge-based coordinate frame', we can define multiple 'hinge-based coordinate frames', each of them centered at a different hinge. In the preprocessing step for each ligand triplet on a single part one should encode the transformations from its 'triplet-based coordinate frame' to all the 'hinge-based coordinate frames' which are located on that part. Thus, e.g., on a part with two hinges, two transformations will be stored for each triplet, while on a part with one hinge only, one transformation will be stored. The recognition step will remain as described

above, except that each receptor triplet will 'vote' for as many frames as the number of different transformations stored in its table entry.

Our discussion focused on hinge-bent ligands, where the two ligand parts are allowed to rotate freely around a point in the 3-D space. However, as was mentioned before, the rotational bonds are more restrictive, since they allow only one degree of freedom rotation around the bond axis. If one knows, that there is only one rotational bound, this restriction can be naturally incorporated in our algorithm and improve its performance. Specifically, in the pre-processing phase one can locate the 'hinge-based coordinate frame' in the middle of the bond and align its x-axis with it. In the recognition stage one would look not only for a high accumulation of votes for a candidate 'hinge-centered coordinate frame' origin, but would also expect the x-axes of these frames to align in order to produce viable hypotheses. This induces a better filtering of the matching algorithm. As mentioned above, in our preliminary implementation we have preferred the more general 'hinge' constraint, since it may well approximate two subsequent rotations as well.

Experimental results

We have implemented the flexible docking algorithm along the lines presented in the previous section. The atom coordinates of the complexed ligand and receptor (considered as an input to our method), were determined by the X-ray crystallography technique and stored in the Protein Data Bank as crystal bound structures (Bernstein *et al.*, 1977). Since the ligand and the receptor were extracted from bound complexes, the 'correct' solutions are those with translations and rotations close to zero. The algorithm was applied to four complexes: the HIV-1 protease complexed with the U-75875 inhibitor; the dihydrofolate reductase (DHFR) complexed with methotrexate, and, separately with NADPH; and lactate dehydrogenase complexed with NAD-lactate. Both rigid body and docking with hinge-bending were carried out. Good results with small root mean square distances (RMSD) as compared with the crystal structure of the complex ('correct' solutions) were obtained.

A plot of verified transformations for docking a rigid NADPH ligand to the DHFR receptor, is presented in Figure 4a. Each diamond corresponds to a transformation represented by its *translation distance* (abscissa) and its *angular distance* (ordinate). The former is the l_2-norm of the transformation's translation vector, and the latter is the evaluation of the term which corresponds to the rotation angle around an equivalent axis ($arccos\frac{tr(R)-1}{2}$, where R is the transformation's rotation matrix and *tr* is its trace). Note that the bottom left cluster of

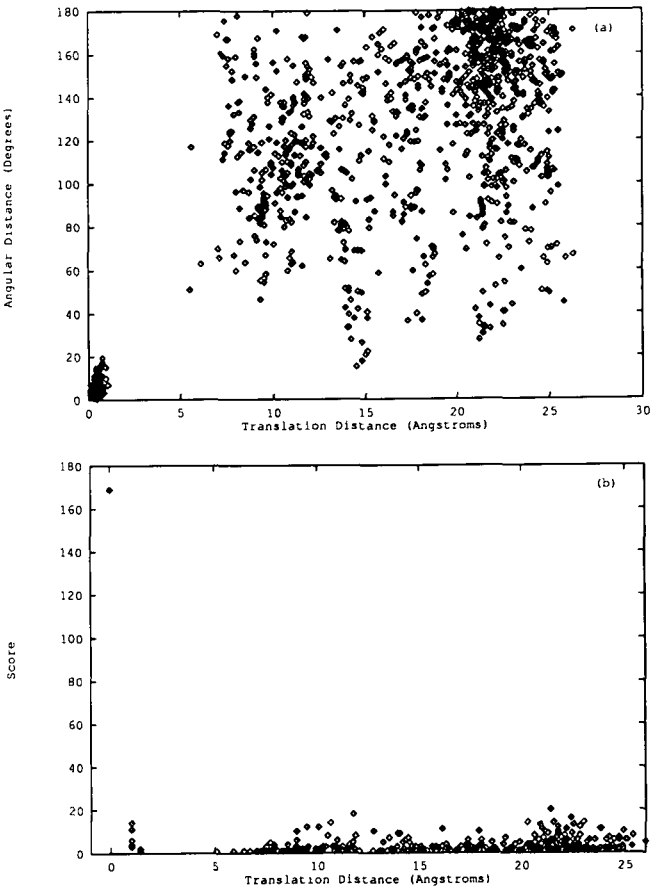

Fig. 4. Rigid docking of the NADPH molecule to a dihydrofolate reductase (DHFR) receptor. (**a**) Geometrically acceptable transformations The translation distance is represented by the X coordinate The Y coordinate represents the angular distance. See text for details. (**b**) The score of transformations shown in (a). The Y coordinate represents the translation distance's score for all rotations. The zero translation distance, having the highest score (above 170), corresponds to the NADPH/DHFR crystal bound transformation.

transformations corresponds to the 'correct' solution. A plot of high scoring solutions for translations over all rotations is shown in Figure 4b. The score (ordinate) is actually the total number of rotations counted for the respective translations (abscissa). The zero translation is the highest scoring one, having rotations with angular distance close to zero degrees (as seen in Figure 4a).

Table 1 summarizes the results obtained for the 'hinge-bending' docking of the aforementioned four complexes. The ligands are assumed to consist of two parts connected by a hinge, as exemplified by Figures 5 and 8. In the NADPH and NAD-lactate ligands, the hinge was located at the O3 atom (see Figure 5 for NADPH). In the methotrexate, the hinge was located at the central carbon atom C9, between the two ring systems (not shown here). In the U-75875 inhibitor, the hinge was located between the two CP atoms (see Figure 8).

The execution times of our program running on the SGI-

Table I. Flexible Docking Parameters and Results

Ligand Receptor	U-75875 HIV-1 protease		NADPH Dihydrofolate reductase		Methotrexate Dihydrofolate reductase		NAD-Lactate Lactate dehydrogenase	
PDB filename	Wlodawer[a]		3DFR		3DFR		5LDH	
Hinge location:	between CP atoms		O3 atom		C9 atom		O3 atom	
Input size:								
ligand's number of interest points	59		48		33		61	
–in first part	36		27		13		33	
–in second part	24		22		21		29	
receptor's number of interest points	351		427		427		621	
Parameters:								
collision distance	1.25 Å		1.25 Å		1.50 Å		1.25 Å	
contact distance	1.00 Å		1.00 Å		1.00 Å		1.00 Å	
voting threshold	20%		20%		20%		20%	
contact threshold	60%		80%		80%		30%	
	Full	Fast[b]	Full	Fast	Full	Fast	Full	Fast
Execution times[c]:								
matching stage	13 7	2 6	21.7	1 9	20.8	1.3	74 5	9 4
collision check[b]	30.5	16.6	31.8	9.8	23.9	3.6	268.4	165.0
self collision check	26.6	1 1	10.3	0 7	2.9	0.6	232 1	26.6
total	70.8	20.3	63.8	12.4	47.6	5.5	575.0	201.0
Best solution:								
contact percentage[d]	(87%)		(98%)		(88%)		(39%)	
	81%	81%	96%	98%	88%	82%	35%	37%
RMSD of first part	0.39 Å	0.39 Å	0 42 Å	0.55 Å	0.60 Å	0.91 Å	0.60 Å	1.05 Å
RMSD of second part	0.60 Å	0.60 Å	0.37 Å	0.37 Å	0.83 Å	1.93 Å	1.62 Å	1.81 Å

[a] Personal communication.
[b] See text for further explanations of the distribution of the times.
[c] The times are given in minutes.
The execution times of the preprocessing step of all cases, is 1 second and less
[d] The percentage appearing in brackets indicates the contact percentage of the original crystal conformation of the ligand.

Iris INDIGO workstation (with 32 MByte of memory) are given in Table 1. The four cases were tested for two types of runs. In the first run (referred to in Table 1 as 'full' for full scan), all triplets of the receptor's interest points were considered for the matching in the recognition step. However, in the second run (noted as 'fast' for fast scan), the receptor's space was divided to eight segments and an adjoining ninth segment. The matching is then conducted for the triplets within each segment. This procedure reduces the total number of triplets handled, and hence, drastically reduces the execution times. We should note that this technique implies that less candidate solutions are obtained, since less triplets are matched, thus the best solution may be affected. A further improvement of the times is introduced in the self collision check. Only binding modes receiving a contact percentage which is higher than the *contact threshold*, are checked for self penetration. In both runs, the actual running times of the collision check procedure are lower when taking into account an implementation problem of the algorithm and its solution (as discussed below). Following the matching stage, we are left with many candidate solutions. It is unfeasible to record all of these for the subsequent verification stage, due to memory (storage) consumption. However, we record the score of the highest scoring hinge location for evaluating the *voting threshold* criterion. The search of the accumulator of votes for obtaining the highest scoring hinge locations, is done by conducting the voting procedure again, but at this stage, recording only the transformations which have passed the *voting threshold* criterion. This solution is time consuming, and since it is done at the initialization step of the collision check procedure, the actual net times of the collision check for the four test-cases, are roughly the difference between the collision check and the matching stage times (as appearing in Table 1). Another possible solution for this problem may be to efficiently record and retrieve the candidate transformations on an external storage device. Additional plausible solution is to monitor the accumulator of votes during the matching stage, and to define an adaptive terminating rule for the voting procedure (see Yla-Jaaski and Kiryati, 1993, for discussion).

For all four cases, successful flexible docking was

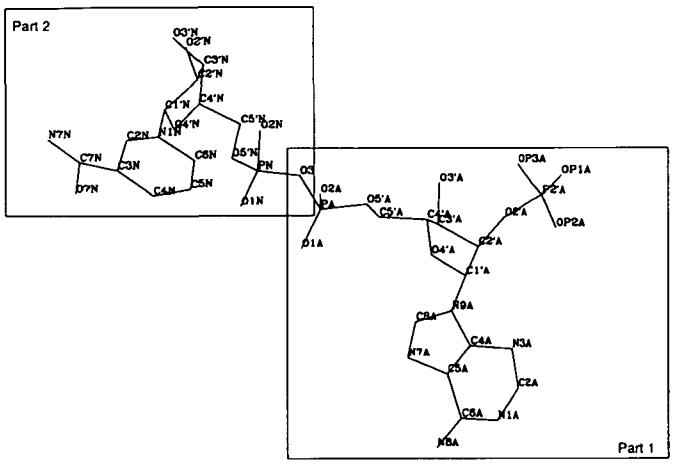

Fig. 5. The spatial structure of NADPH, showing the two parts of the ligand joined by the hinge located at the O3 atom.

achieved, with solutions having relatively small RMSD (see Table 1). The figures presented here are the results obtained for the NADPH (Figures 6, 7) and U-75875 inhibitor (Figures 9, 10) flexible docking. Consistent results have been obtained for the other two complexes (not shown here). The transformations presented and discussed below are between the 'hinge-based coordinate frame' and the 'candidate hinge-centered reference frame'. The former is in the ligand, and the latter is in the negative image of the receptor (see Figure 3). Figures 6a and 6b show the angular distance as a function of the translation distance of the generated transformations, for each of the two parts separately. The transformations obtained are the ones passing the verification stage. The ligand and the receptor have been extracted from the complex, without rotating or translating them. The 'correct' solutions, therefore, are those with rotations and translations near

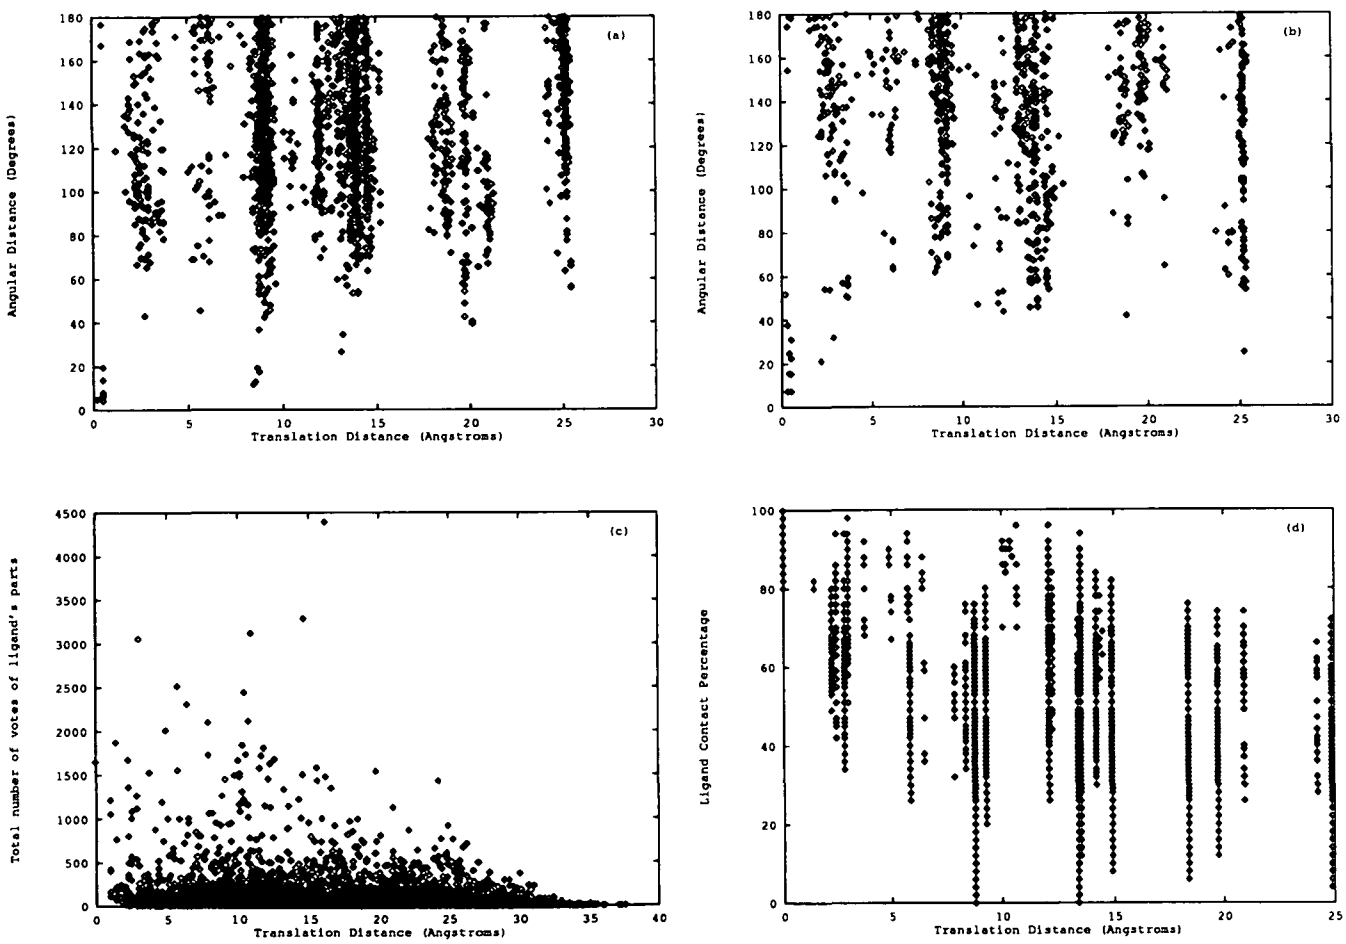

Fig. 6. The NADPH/DHFR flexible docking. (a) Geometrically acceptable transformations of the first part of the NADPH ligand (b) Geometrically acceptable transformations of the seconds part of the NADPH ligand. The transformations are represented by their 'angular distance' (ordinate) as a function of their 'translation distance' (abscissa). See text for details. (c) The number of votes as a function of the translation distance received by the two parts of the NADPH 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. (d) The 'goodness' of the generated binding modes. The percentage of the NADPH ligand's van der Waals spheres which are in contact with the DHFR receptor spheres (ordinate), are plotted as a function of translation distances for the different binding modes

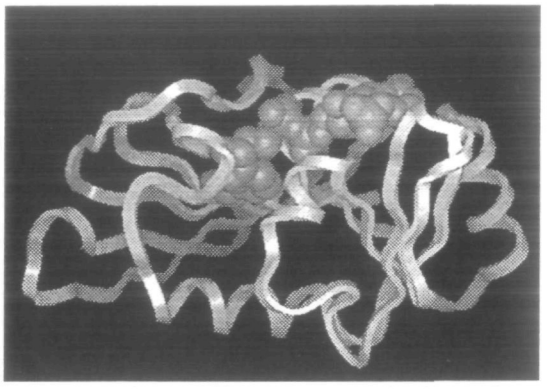

(a)

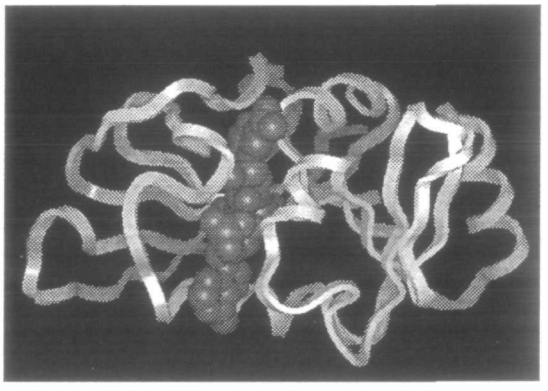

(b)

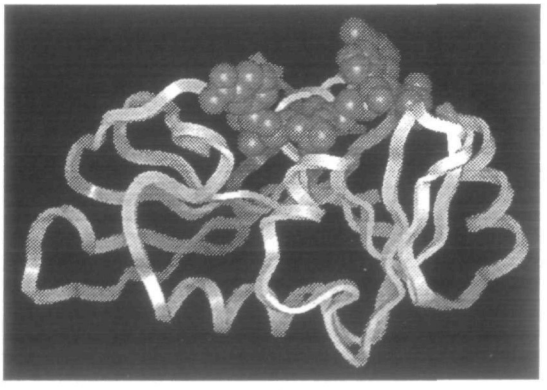

(c)

Fig. 7. Examples of the NADPH/DHFR flexible docking. The ligand molecule is drawn in its van der Waals solid sphere representation. The DHFR molecule is represented as a ribbon. The location and orientation of the DHFR is the same in all figures. (**a**) The original crystal bound complex. The contact percentage of the ligand is 98%. (**b**) A predicted binding mode of the NADPH ligand having a contact percentage of 92% The translation distance of the transformation is 10.4 Å. The angular distances are 127.3 and 155.7 degrees for the first and the second parts, respective to the parts of the crystal bound conformation of the ligand. (**c**) An additional predicted binding mode of the NADPH ligand having a contact percentage of 94%. The translation distance of the transformation is 2.8 Å. The angular distances are 65.5 and 158.3 degrees for the first and the second parts, respective to the parts of the crystal bound conformation of the ligand.

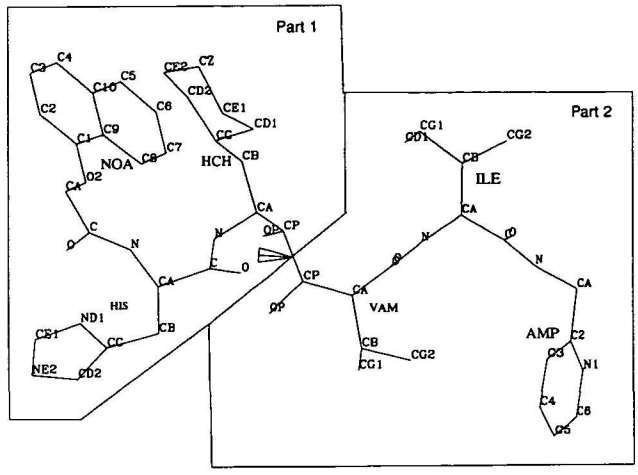

Fig. 8. The spatial structure of the U-75875 inhibitor, showing the two parts of the ligand, joined by the hinge pointed at by the arrow. The hinge is located between the CP atoms.

zero. Both plots depict that such a 'cluster' of solutions is obtained. The transformations obtained for the inhibitor docking are shown in Figures 9a and 9b.

Figures 4a, 6ab and 9ab are 2-D diagrams which theoretically should have been drawn in six dimensions. Each transformation in space, is actually represented by the three dimensions of the translation (the x, y and z coordinates) and the three dimensions of the rotation (the three rotation angles). Since drawing a 6-D diagram is infeasible, the two dimensions representing the transformations are the translation and angular distances. These distances, implicitly represent the transformations (the six aforementioned parameters), by displaying the generated solutions qualitatively. One should note for example, that as the translation distance grows, the number of possibilities to obtain it by different x, y and z also grows. This stems from the fact that three dimensions are collapsed along each of the two axes. Thus, the farther we are from the origin, the more we encounter different solutions collapsing there at random. This creates a visual illusion of peaks and clusters of solutions which are far from the 'zero translation' in addition to the 'correct', close to zero solutions. In the real 6-D parameter space, these high peaks are spread among many low scoring candidate solutions.

The matching votes of the NADPH two parts combined, (prior to checking the transformations score for the voting threshold criterion), are shown in Figure 6c. The translations which are close to zero, i.e., the 'correct' translations, are among the twenty high scoring ones. Some of the high scoring translations, are later rejected by the collision and self-collision checks, leaving a cluster of 'correct' transformations as shown previously in Figures 6a, 6b. The matching votes of the inhibitor docking are

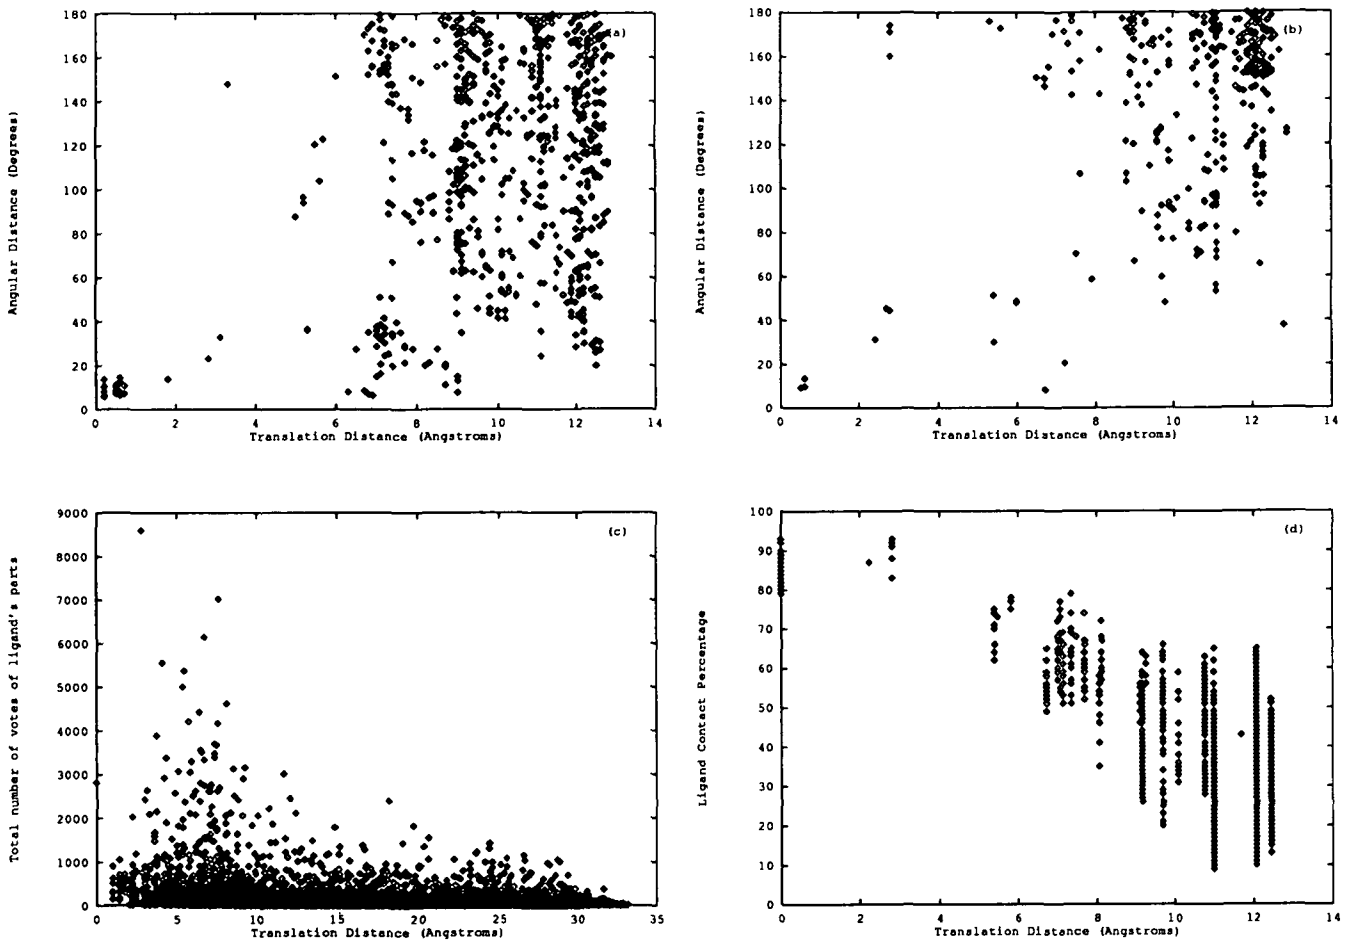

Fig. 9. The U-75875 inhibitor/HIV-1 protease flexible docking. (**a**) Geometrically acceptable transformations of the first part of the inhibitor (consisting of the NOA, HIS and HCH groups). (**b**) Geometrically acceptable transformations of the seconds part of the inhibitor (consisting of the VAM, ILE and AMP groups) The transformations are represented by their 'angular distance' (ordinate) as a function of their 'translation distance' (abscissa). See text for details. (**c**) The number of votes as a function of the translation distance received by the two parts of the inhibitor. 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. (**d**) The 'goodness' of the generated binding modes. The percentage of the inhibitor's van der Waals spheres which are in contact with the receptor spheres (ordinate), are plotted as a function of translation distances for the different binding modes.

shown in Figure 9c. The 'correct' translations are among the twenty seven high scoring ones.

The 'goodness' of the solutions is shown in Figure 6d. This figure depicts the percentages of NADPH spheres which are in contact with the DHFR receptor spheres as a function of the translation distance. As the figure shows, the largest contact-percentages are obtained at near zero translation distances, yielding the reconstituted crystal complex. (See also the inhibitor's results in Figure 9d). However, some additional binding modes yield a high contact percentage as well. These may constitute alternate binding sites of the ligand to the receptor. The high scoring alternate binding modes of the ligand were investigated. Although the contact percentage scored by NADPH/DHFR the crystallographic complex is 98% (see Figure 7a), additional high scoring binding modes were also generated. We chose to present other relatively high scoring binding modes, having contact percentage of 92% and 94% (Figures 7b and 7c, respectively). In comparison with the original complex, the inhibitor is docked in other, alterate, locations of the receptor, with significant changes in the hinge locations and orientations of the parts. Geometrically good-fitting predictive binding modes have been obtained for the inhibitor in Figures 10b and 10c.

Discussion and future research

Here we present a general approach, of docking a ligand onto a receptor, allowing hinge-bending motions of relatively rigid parts. Our main algorithmic tool is the generalization and the extension of the Geometric Hashing and Generalized Hough Transform techniques, which were originally developed for partially occluded

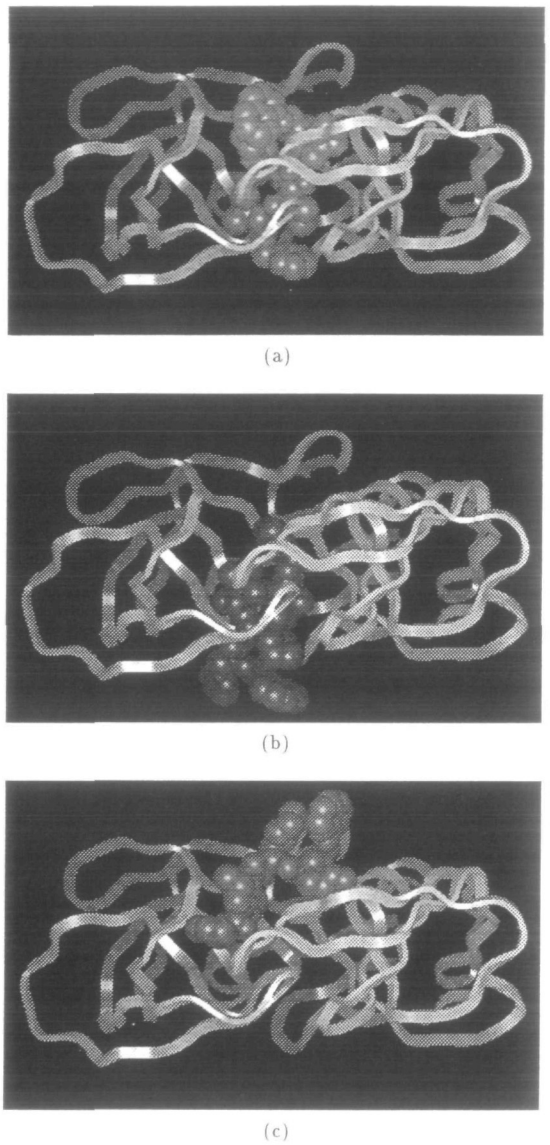

Fig. 10. Examples of the U-75875 inhibitor/HIV-1 protease flexible docking. The inhibitor molecule is drawn in its van der Waals solid sphere representation. The HIV-1 protease molecule is represented as a ribbon. The location and orientation of the HIV-1 protease is the same in all figures. (**a**) The original crystal bound complex. The contact percentage of the inhibitor is 87%. (**b**) A predicted binding mode of the inhibitor having a contact percentage of 79%. The translation distance of the transformation is 7.3 Å. The angular distances are 155.2 and 150.1 degrees for the first and the second parts, respective to the parts of the crystal bound conformation of the inhibitor. (**c**) An additional predicted binding mode of the inhibitor having a contact percentage of 62%. The translation distance of the transformation is 11 Å. The angular distances are 69 and 97.7 degrees for the first and the second parts, respective to the parts of the crystal bound conformation of the inhibitor.

articulated object recognition in Computer Vision & Robotics. We apply it to the case of docking of a ligand into a large receptor, obtaining the 'correct' crystal bound complex and additional alterate binding modes. The analogy between the type of problems encountered in Computer Vision and in Molecular Biology has brought about this interdisciplinary research endeavour (Nussinov and Wolfson, 1991).

The method described here is geometrical in nature, searching for optimal fit of patches of surfaces between the receptor and the ligand, allowing hinge-bending. Potential solutions resulting in collisions between ligand and receptor atoms and self collision between the parts are discarded. Following these procedures, we are still often left with many geometrical solutions. Geometrically acceptable docked solutions should be the input to routines examining the chemical interactions between the receptor-ligand atoms at the interface. Calculation of the energies of the docked structures will enable discriminating between the biologically favorable and unfavorable solutions. An efficient routine for energetical evaluation of the docked structures is currently being developed (S.L.Lin, personal communication).

The range of the potential implications of our approach, encompasses investigations in a number of directions. First among these is the positioning of the hinge-bending movement in the protein-receptor. Second, one would like to allow motions at some hinges during the recognition and binding of two (relatively large) protein molecules. Here, we have treated the case of the docking of a small ligand into a protein receptor. Third, whereas the algorithm presented here is for the simpler, single hinge case, it can be extended to handle multiple hinges (Wolfson, 1991). Motions can be allowed simultaneously on all hinges. An automated identification of flexible hinges (Segawa and Richards, 1988) would further automate this procedure. These tasks are facilitated by the recently developed surface representation which is highly suitable to our purposes (Lin *et al.*, 1994). This representation describes the surface in terms of accurately—and sparsely—placed points and their surface normals, representing the 'caps' (convex), 'pits' (concave) and 'belts' (saddle) regions of the surface. This representation is robust, and is independent of the initial density of dots used to describe the surface (Connolly, 1983b; Connolly, 1983a). Extensive testing with rigid-body docking has indicated the superiority of this surface representation to other surface routines that we have tried (Fischer *et al.*, in preparation).

The implications of this Computer Vision & Robotics-based tool are not, however, restricted to docking. One can easily envisage a much wider range, where such an algorithm can yield a new insight into the understanding of protein structures, into evolutionary implications and as such, could aid in aspects of protein folding. Recently a Computer Vision based algorithm has been adapted for the search of 3-D, sequence-order independent, substructural motifs in protein molecules (Bachar *et al.*,

1993; Fischer *et al.*, 1992). By allowing motions on several hingess, new 3-D, substructural motifs could be detected in protein structures. Thus the algorithm's application to comparisons and investigations of protein structures is likely to yield 3-D, sequence-order independent, unpredefined, hinge-bent motifs.

Our implemented approach for the recognition of rotated and translated 3-D molecular structures with rotary hinges (joints), potentially provides an implementation for solving 3-D to 3-D matching problems in Computer Vision. These problems arise from fields such as medical imaging, CAD/CAM, military and space systems, robotics assembly, etc.

Such diverse investigations are enabled by our novel, Computer Vision & Robotics based algorithm, along the degrees of freedom it affords.

Acknowledgments

We would like to thank A. Wlodawer for the coordinates of the HIV-1 protease–U-75875 inhibitor complex. 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 Program Resources, Inc. 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 has been supported in part by a grant from the *Israel Science Foundation* administered by the *Israel Academy of Sciences*, and by grant No. 83-00481 from the U.S.-Israel Binational Science Foundation (BSF), Jerusalem, Israel. The research of R. Nussinov in Israel has been supported in part by grant No. 91-00219 from the BSF, and by a grant from the *Israel Science Foundation* administered by the *Israel Academy of Sciences*. This work formed part of the Ph.D. Thesis of B. Sandak, Tel Aviv University.

References

Bachar,O., Fischer,D., Nussinov,R., and Wolfson,H.J. (1993). A computer vision based technique for 3-D sequence-independent structural comparison of proteins *Prot. Eng.* **6**, 279–289.

Ballard, D. (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, **13(2)**, 111–122

Beinglass,A. and Wolfson,H J (1991). Articulated object recognition, or, how to generalize the generalized Hough transform. In: *Proceedings of the IEEE Computer Vision and Pattern Recognition Conf.* pp. 461–466: Maui, Hawaii.

Bennet,W. and Haber,R. (1984). Structural and functional aspects of domain motion in proteins. *CRC Crit. Rev. Biochem*, **15**, 291.

Bernstein,F., Koetzle,T., Williams,G., Meyer,E., M.D., B., Rodgers,J., Kennard,O., Shimanonchi,T., and Tasumi,M. (1977). The protein data bank: A computer-based archival file for macromolecular structures. *J Mol. Biol.* **112**, 535–542.

Brooks,R. (1981). Symbolic reasoning around 3-D models and 2-D images. *Artificial Intelligence*, **17**, 285–348.

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., 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.

Dixon,M., Nicholsen,N., Sherwchuk,L., Baase,W., and Matthews,B. (1992). Structure of a hinge-bending Bacteriophage T4 Lysozyme Mutant, Ile3 → Pro. *J. Mol Biol* **227**, 917–933.

Faber,H and Matthews,B. (1990). A mutant T_4 lysozyme displays five different crystal conformations *Nature*, **348**, 263–266.

Fischer,D., Norel,R., Nussinov,R., and Wolfson,H.J. (1993) 3-D docking of protein molecules. In: *Fourth Symposium on Combinatorial Pattern Matching* pp. 20–34, Padova, Italy. Springer Verlag. Lecture Notes in Computer Science 684.

Fischer,D., Nussinov,R., and Wolfson,H.J. (1992) 3-D substructure matching in protein molecules. In: *Third Symposium on Combinatorial Pattern Matching* pp. 136–150, Tucson, Arizona, USA· Springer Verlag Lecture Notes in Computer Science 644.

Gerstein,M. and Chothia,C. (1991). Analysis of protein loop closure. Two types of hinges produce one motion in Lactate Dehydrogenase *J. Mol. Biol.* **220**, 133–149.

Gerstein,M., Lesk,A., and Chotia,C. (1994). Structural mechanisms for domain movements in proteins. *Biochemistry*, **33**, 6739–6749.

Ghose,A. and Crippen,G. (1985). Geometrically feasible binding modes of a flexible ligand moledule 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.

Grimson,W. (1991). Recognition of object families using parametrized models In: *Proceedings of the IEEE Computer Vision and Pattern Recognition Conf.* pp. 566–572: Maui, Hawaii.

Jiang,F. and Kim,S. (1991). Soft docking: Matching of molecular surface cubes. *J. Mol. Biol.* **219**, 79–102.

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. on Robotics and Automation*, **6**, 578–589.

Leach,A. and Kuntz,I. (1992). Conformational analysis of flexible ligands in macromolecular receptor sites. *J. Comp. Chem.* **13**, 730–748

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

Lowe,D.G. (1991). Fitting Parametrized Three-Dimensional Models to Images. *IEEE Trans. on Pattern Analysis and MAchine Intelligence*. **13**, 441–450

Miller,M., Schneider,J , Sathyanarayana,B , Toth,M , G R , M., Clawson,L , Selk,L., Kent,S , and Wlodawer,A. (1989). Structure of complex of synthetic HIV-1 protease with a substrate-based inhibitor at 2.3 Å resolution *Science*. **246**, 1149–1152.

Nussinov,R. and Wolfson,H.J. (1991). Efficient detection of three-dimensional motifs in biological macromolecules by computer vision techniques. *Proc. Natl. Acad. Sci. USA*. **88**, 10495–10499.

Rini,J., Schulze-Gahmen,U , and Wilson,I. (1992). Structural evidence for induced fit as a mechanism for antibody-antigen recognition. *Science*, **255**, 959–965.

Sali,A., Veerapandian,B., Cooper,J.B., Moss,D., Hofmann,T., and Blundell,T. (1992). Domain flexibility in aspartic proteinases *Proteins· Struct., Funct., Genet.* **12**, 158–170.

Segawa,S. and Richards,F. (1988). Identification of regions of potential

flexibility in protein structures: Folding units and correlations with intron positions. *Biopolymers*, **27**, 23–40.
Shakunaga,T (1991). Pose estimation of jointed structures. In. *Proceedings of the IEEE Computer Vision and Pattern Recognition Conf.* pp 566–572: Maui, Hawaii.
Stanfield,R., Fieser,T., Lerner,R., and Wilson,I. (1990). Crystal structure of an antibody to a peptide and its complex with peptide antigen at 2.8 Å. *Science*, **248**, 712–719.
Wang,H. (1991). Grid-search molecular accessible algorithm for solving the protein docking problem. *J. Comp. Chem.* **12**, 746–750.
Weber,C., Wilder,G., von Freyberg,B., Traber,R., Braun,W., Widmer,H , and Wuthrich,K. (1991). The NMR structure of cyclosporin A bound to cyclophilin in aqueous solution. *Biochemistry*, **30**, 6563–6574.
Wolfson,H.J. (1991). Generalizing the generalized Hough transform *Pattern Recognition Letters*, **12**, 565–573.
Yla-Jaaski,A and Kiryati,N. (1993) Adaptive termination of voting in probabilistic Hough algorithms. In: *The 10th Israeli Symposium on Artificial Intelligence, Computer Vision and Neural Networks* pp. 315–324, Ramat-Gan, Israel: IPA.

Presented at the IEEE Conference on June 20, 1994; accepted on November 15, 1994