Docking Flexible Ligands to Macromolecular Receptors by Molecular Shape

Renee L. DesJarlais,[†] Robert P. Sheridan,*[‡] J. Scott Dixon,[‡,§] Irwin D. Kuntz,*[†] and R. Venkataraghavan[‡]

Department of Pharmaceutical Chemistry, University of California at San Francisco, San Francisco, California 94143, and Lederle Laboratories, Pearl River, New York 10965. *Received January 21, 1986*

We present a method to explore the interaction of flexible ligands with receptors of known geometry on the basis of molecular shape. This method is an extension of that described by Kuntz et al. (*J. Mol. Biol.* **1982**, *161*, 269). The shape of a binding site on a macromolecular receptor is represented as a set of overlapping spheres. Each ligand is divided into a small set of large rigid fragments that are docked separately into the binding site and then rejoined later in the calculation. The division of ligands into separate fragments allows a degree of flexibility at the position that joins them. The rejoined fragments are then energy minimized in the receptor site. We illustrate the method with two test cases: dihydrofolate reductase/methotrexate and prealbumin/thyroxine. For each test case, the method finds binding geometries for the ligand near that observed crystallographically as well as others that provide good steric fit with the receptor.

The interaction of macromolecular receptors and their small molecule ligands is an essential step in many biological processes: regulatory mechanisms, the pharmacological action of drugs, the toxic effect of certain chemicals, etc. The ability to predict such interactions leads to the possibility of designing a compound directed at a specific target; such a compound could prove useful as a research tool or a novel drug.

When detailed structural information about the receptor is not available, a model must be deduced from the ligands that bind to it. A number of statistical methods are commonly used in this case. These include Hansch analysis[1] and distance geometry methods[2] that predict characteristics of a site by relating the selected structural properties of active compounds to their biological activities. While these approaches avoid the need for detailed knowledge of the receptor, which is unavailable for most biological systems, they require much experimental data on many compounds and provide no straightforward way to predict the binding of compounds that are very different from the compounds used to derive the data. In contrast, it is also possible to start from a receptor site of known structure.[3] This usually requires an X-ray structure of the receptor or a closely related molecule that can serve as a model for the receptor.

Early attempts to predict molecular interactions at receptor sites of known geometry used hand-constructed models to visualize the receptor site and chemical intuition to determine the final bound geometry of the ligand.[3] Recent developments in computer graphics have made the task of molecular modeling much simpler and faster. The determination of the bound geometry of a ligand, however, is still mainly a manual process that depends critically on the chemist's intuition. Such determinations have been successful in several cases.[4,5] New techniques for displaying the chemical environment of a receptor site as a potential energy grid[6] can help guide the chemist in modeling the most likely bound conformation. While all of these developments aid the process of manually fitting a ligand to its receptor, an automatic method for generating reasonable binding modes is desirable. Such a search could suggest modes that might not be obvious to the chemist. Determination of whether each suggested mode is chemically reasonable could be made by visual inspection or by the evaluation of an energy function.

There have been several attempts to automate the process of fitting ligands to their receptors. Wodak and Janin[7] studied the interaction of two proteins, using a simplified energy representation of the molecular structure. They searched four of the six degrees of freedom between the molecules to find a possible binding modes. Santavy and Kypr[8] have also developed a geometry-based method to search for the optimum orientation of two macromolecules. Kuntz et al.[9] presented an algorithm designed to fit small molecules into their macromolecular receptors. The last two methods are concerned with complementarity of shape and do not attempt to evaluate an energy function at the binding site. Both methods use the solvent accessible surface as a starting place from which to characterize shape, but they differ in how shape is represented as well as in how shapes are matched. All the methods discussed so far are restricted to rigid ligands and receptors. In real systems, both the ligand and receptor have some degree of flexibility. In this paper, we describe a new approach to the study of flexible ligands wherein a ligand is approximated as a small set of rigid fragments. The algorithm of Kuntz et al.[9] is used to find a set of bound geometries for each fragment. These geometries are then scanned for arrangements such that the fragments can be rejoined to recreate the ligand, which is then energy minimized on the receptor. We present the results of two test cases: binding of methotrexate to dihydrofolate reductase and thyroxine to prealbumin.

Methods

Our goal is to find a series of sterically reasonable binding geometries for a given ligand in a receptor binding site whose structure is known. Our algorithm can be summarized:

(1) Generate a set of spheres representing the shape of the receptor binding site from the macromolecular surface.

(2) Approximate the ligand as a small number of large rigid fragments.

(3) For each fragment, match fragment atom–atom distances with receptor sphere–sphere distances to find sets

[†] University of California at San Francisco.

[‡] Lederle Laboratories.

[§] Present address: Smith Kline French Laboratories, Philadelphia, PA 19101.

(1) Hansch, C.; Leo, A. *Substituent Constants for Correlation Analysis in Chemistry and Biology*; Wiley: New York, 1979.
(2) Crippen, G. M. *J. Med. Chem.* **1979**, *22*, 988.
(3) Beddell, C. R.; Goodford, P. J.; Norrington, F. E.; Wilkinson, S.; Wootton, R. *Br. J. Pharmacol.* **1976**, *57*, 201.
(4) Blaney, J. M.; Jorgensen, E. C.; Connolly, M. L.; Ferrin, T. E.; Langridge, R.; Oatley, S. J.; Burridge, J. M.; Blake, C. C. F. *J. Med. Chem.* **1982**, *25*, 785.
(5) Freudenreich, C.; Samama, J. P.; Biellmann, J. F. *J. Am. Chem. Soc.* **1984**, *106*, 3344.
(6) Goodford, P. J. *J. Med. Chem.* **1985**, *28*, 849.
(7) Wodak, S. J.; Janin, J. *J. Mol. Biol.* **1978**, *124*, 323.
(8) Santavy, M.; Kypr, J. *J. Mol. Graphics* **1984**, *2*, 47.
(9) Kuntz, I. D.; Blaney, J. M.; Oatley, S. J.; Langridge, R.; Ferrin, T. E. *J. Mol. Biol.* **1982**, *161*, 269.

0022-2623/86/1829-2149$01.50/0

of receptor spheres that can accommodate the fragment shape. Orient the fragment in the receptor site.

(4) For each fragment, eliminate those orientations that result in a significant overlap between the fragment and the receptor, and eliminate those orientations which are redundant.

(5) To recreate the ligand, systematically pair the orientations of two fragments. Save those orientations in which specified atoms from each fragment are close enough that the fragments can be rejoined as they were joined in the intact ligand. This step is repeated, if necessary, until all fragments are joined and the entire ligand is regenerated. A set of orientations for each fragment chosen in this way constitutes a binding "mode" for the ligand to the receptor.

(6) Divide the binding modes into families and energy minimize each family in the presence of the receptor.

We now discuss these steps in detail, with emphasis on the steps that are different from those in the original algorithm of Kuntz et al.[9]

Characterization of Shape. The characterization step begins with a calculation of the solvent accessible surface as described by Richards[10] using a program developed by Connolly.[11-13] The surface consists of a series of points representing the concave (reentrant) or convex (contact) features. Conventionally, a probe radius of 1.4 Å is used. For our purpose, surfaces are calculated using only nonhydrogen atoms with appropriate "united atom" radii.

We generate from the surface points a set of spheres representing the negative image of the receptor volume. (In the "lock-and-key" analogy the receptor spheres represent the shape of the keyhole.) Each sphere is characterized by its radius and by the location of its center, but we use only the latter information. This algorithm has been previously described.[9] For this study, we have eliminated the previous requirement that surface points taken pairwise in the sphere-generating algorithm must be on atoms at least four amino acid residues apart in sequence. Once generated, the spheres are separated into "clusters", in which each sphere in a cluster overlaps at least one other sphere in the cluster. Over an entire macromolecule there are usually several distinct clusters, representing surface invaginations of various sizes. While each site could be studied in turn,[9] we focus on the cluster with the most spheres. In most macromolecules, including those in our test cases, that cluster corresponds to the crystallographically determined binding site.

The ligand is separated into the largest fragments that can be treated as being rigid. An atom may appear in more than one fragment. Previously, the shape of a ligand was characterized by spheres in the same way as the receptor, except that the reentrant surface points were used, generating a positive image of the volume. For this study, we use atom centers to define the shape of each fragment. This simplification is appropriate whenever the fragments are small.

Matching and Orientation of Fragments. Having characterized the shape of the ligand fragment and the receptor binding site, we find the geometrically possible ways to orient each fragment in the receptor site using systematic distance matching.[9] Each fragment atom i is systematically paired with a receptor sphere center k. A second pair, atom j with sphere l, is accepted if the distances obey the condition:

$$abs(d(i,j) - d(k,l)) \leq C$$

where C is a parameter set by the user. Additional pairs are assigned until no further pairs meet the condition. The minimum number, N, of atom-sphere pairs needed for a "match" to be saved can be set by the user, although at least four pairs are necessary to determine a unique docking. Because we are dealing with fragments that are much smaller than the receptor site, we have eliminated the additional requirement in Kuntz et al.[9] that the center of gravity of the ligand be nearly coincident with the center of the receptor spheres.

For each match we use the least-squares algorithm of Ferro and Hermans[14] to obtain a rotation/translation matrix that will best superimpose each fragment atom onto its paired receptor sphere. The calculated matrix is then applied to the whole fragment to generate an "orientation", i.e., coordinates for the fragment positioned in the receptor binding site. There are generally on the order of n times m matches for each fragment, where n is the number of receptor spheres and m is the number of fragment atoms.

Filtering the Orientations. Here we take a different approach from that of Kuntz et al.[9] who attempted to remove overlap of a ligand and receptor atoms by displacing a particular orientation by small increments. This was the most computationally intensive step, but it was not particularly effective either in improving the overlap or filtering out structures that could not be improved. We have chosen to simply discard any orientation in which a fragment atom is within 2.5 Å of a receptor atom. This has the effect of removing from consideration all orientations that have significant overlaps with the receptor.

Many matches are degenerate, i.e., they result in orientations that are nearly coincident. We can remove redundant orientations by calculating the root mean square deviation in atomic position (rms) for every pair of orientations and discarding all but one member of a group within which the rms is very small.

Joining the Fragments and Refining the Ligand. Once each fragment of the ligand has been separately matched, oriented, and filtered, the fragments must be joined to regenerate the full ligand. The joining step systematically pairs orientations of two fragments. The user specifies a maximum and minimum distance between an atom from one fragment and an atom from another fragment. Any number of distances can be specified. Pairs of orientations that meet the distance criteria and that have no other atoms within 2.5 Å of each other (i.e., no significant overlap between fragments) are saved. If a third fragment is to be joined, the joining step is repeated, this time with the orientations of the third fragment systematically paired with saved pairs from the first joining step. Fragments can be added sequentially in this way until the full ligand is regenerated.

Each set of joined fragment orientations saved represents a receptor-binding mode for that ligand. Modes can be compared by an rms deviation in the same way as orientations of a specific fragment are compared. We find it convenient to group modes that bind to the receptor in basically the same way into "families" on the basis of a less stringent rms criterion.

Representative members of each family are energy minimized in the binding site by using the molecular mechanics program AMBER,[15] with the new force field parameters described by Weiner et al.[16] Energy minimiza-

(10) Richards, F. M. *Annu. Rev. Biophys. Bioeng.* **1977**, *6*, 151.
(11) Connolly, M. L. Thesis, University of California, 1981.
(12) Connolly, M. L. *J. Appl. Crystallogr.* **1983**, *16*, 548.
(13) Connolly, M. L. *Science (Washington, D.C.)* **1983**, *221*, 709.
(14) Ferro, D. R.; Hermans, J. *Acta Crystallogr., Sect. A* **1977**, *A33*, 345.
(15) Weiner, P. K.; Kollman, P. A. *J. Comput. Chem.* **1981**, *2*, 287.

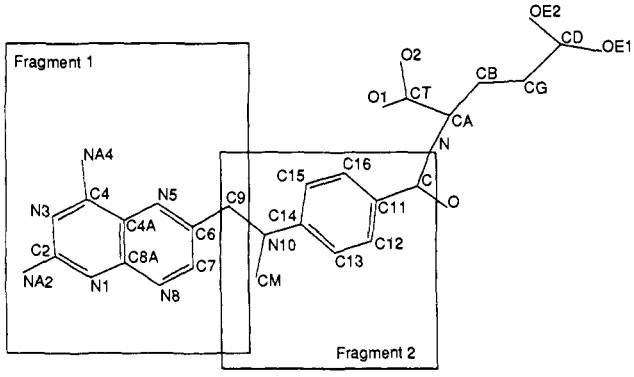

Figure 1. The methotrexate molecule showing division into fragment 1 and fragment 2.

tion is important for several reasons. First, the bonds and angles at the joint between two fragments are likely to be strained, especially if the distance criteria in the joining step are permissive. It is desirable to relieve that strain. Also, minimization allows a degree of induced fit between the ligand and the receptor. Finally, the minimized energies for the various families can be compared.

Results

We have explored two test cases in detail: the binding of methotrexate (MTX) to dihydrofolate reductase from *Lactobacillus casei* (DHFR) and the binding of thyroxine (T4) to human prealbumin (PAB). The structure of both complexes is known from experimental data and model-building. The cases offer an opportunity to test the ability of our method to reproduce a known geometry as well as finding novel geometries that might be of interest in designing new compounds.

Dihydrofolate Reductase/Methotrexate. DHFR is an enzyme of particular interest because some of its inhibitors are useful as antibiotics and chemotherapeutic agents. The structure of *Lactobacillus casei* DHFR complexed with NADPH and MTX is known from X-ray crystallography to 1.7 Å.[17] The set of coordinates we used here is called 3DFR in the Brookhaven Protein Data Bank.[18] None of the X-ray waters was included in the docking.

The structure of MTX is shown in Figure 1. For the purpose of this study, MTX was approximated as two rigid fragments, called fragment 1 and fragment 2 in the figure. The flexible L-glutamate portion of the molecule was neglected, but useful results are obtained despite this omission. Hereafter "ligand" refers to that portion of MTX accounted for by the two rigid fragments.

The results of each step in the match are outlined in Table I. The ligand was regenerated by requiring that atom C9 from each fragment be within 1 Å of each other, since they are the same atom in the intact ligand, and that the atoms N10 and C6 be 2–3 Å apart, so that N10–C9–C6 form a reasonable angle.

The 14 fragment pairs, each corresponding to a "ligand mode", were examined and divided into four families on the basis of their rms deviations from each other. A representative of each family was energy minimized using

Table I. Binding Modes for Dihydrofolate Reductase/Methotrexate

	fragment 1	fragment 2
atoms	13	10
initial matches[a]	432	175
orientations with no overlaps[b]	137	146
nonredundant orientations[c]	122	117
joined solutions (modes)[d]	14	
families[e]	4	

[a] We required a minimum of five atom–sphere pairs for a match with the distance error parameter C = 1.5 Å. [b] An overlap is defined as having any fragment atom within 2.5 Å of any receptor atom. [c] An orientation is considered redundant if it is within 0.1-Å rms deviation of any other orientation. [d] A joined solution was saved if the following distance conditions were met: C9 fragment 1 0.0–1.0 Å from C9 fragment 2; C6 fragment 1 2.0–3.0 Å from N10 fragment 2. All atoms not specified above in either fragment must be more than 2.5 Å from all atoms in the other fragment. [e] Modes are in the same family if they differ by less than 1.0 Å, rms.

Table II. Energies of Binding Modes for Dihydrofolate Reductase/Methotrexate[a]

	rel energy, kcal/mol		rel energy, kcal/mol
x-ray crystal structure	0.0	family 3	−6.0
family 1	80.0	family 4	−3.0
family 2	94.0		

[a] Energies were calculated with AMBER[15] using the new force field parameters.[16] Only those atoms accounted for by fragments 1 and 2 were included for the ligand and only those residues in dihydrofolate reductase within 8 Å of the X-ray position of methotrexate[17] were included in the calculation. The absolute energy of the X-ray position of the ligand after minimization was −266.2 kcal/mol.

AMBER.[15] The calculation included the ligand and all DHFR residues that had at least one atom within 8 Å of the X-ray crystal structure of MTX. Where the polypeptide chain of DHFR was cut, the terminal N and C atoms were constrained to their original positions with a force constant of 100 kcal/(mol Å^2). Before minimization, explicit hydrogens were added to the heteroatoms in the ligand and to the heteroatoms in residues from DHFR. The N1 of MTX was considered to be protonated.

The relative energies obtained after minimization are listed in Table II. In all cases the protein minimized to the same conformation with the exception of a few DHFR residues whose position is shifted slightly depending on the position of the ligand. The overall position of the ligand is changed little by energy minimization, but the bond lengths and angle at the joint between fragments are corrected.

The four families of "ligand modes" fall into two classes on the basis of energy: one group with energies comparable to that of the minimized X-ray structure (Figure 2), and one with the energies considerably higher. The two families of higher energy (1 and 2) could be called "inverted" modes, in that fragment 2 of the ligand is fit into the subsite that holds fragment 1 in the X-ray structure and fragment 1 is in the subsite that holds fragment 2. It is interesting to note that these inverted modes leave no room for the glutamine portion of MTX which was omitted from the ligand. Family 4 is very similar to the X-ray mode. Family 3, however, has the pteridine ring rotated 180° with respect to the X-ray mode. This last mode corresponds to the known binding geometry of folate.[19]

Prealbumin/Thyroxine. The hormone T4 binds to the tetrameric form of PAB in a twofold symmetric cleft formed between two monomers. The crystal structure of

(16) Weiner, S. J.; Kollman, P. A.; Case, D. A.; Singh, U. C.; Ghio, C.; Alagona, G.; Profeta, S., Jr.; Weiner, P. *J. Am. Chem. Soc.* **1984**, *106*, 765.
(17) Bolin, J. T.; Filman, D. J.; Matthews, D. A.; Hamlin, R. C.; Kraut, J. *J. Biol. Chem.* **1982**, *257*, 13650.
(18) Bernstein, F. C.; Koetzle, T. F.; Williams, G. J. B.; Meyer, E. F., Jr.; Brice, M. D.; Rodgers, J. R.; Kennard, O.; Shimanouchi, T.; Tasumi, M. *J. Mol. Biol.* **1977**, *112*, 535.
(19) Filman, D. J.; Bolin, J. T.; Matthews, D. A.; Kraut, J. *J. Biol. Chem.* **1982**, *257*, 13663.

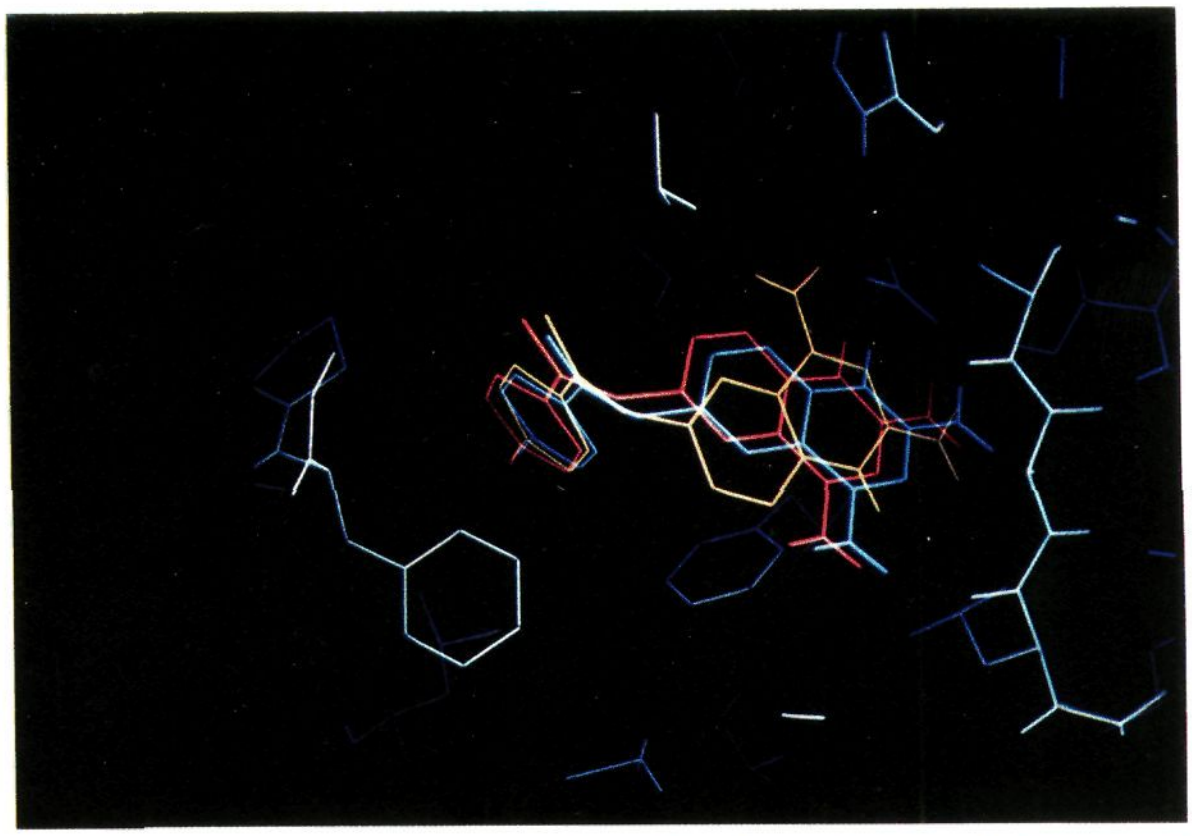

Figure 2. Selected methotrexate binding modes after energy minimization: X-ray[17] (blue), family 3 (red), and family 4 (yellow). The structure of dihydrofolate reductase is that minimized in the presence of methotrexate bound in the X-ray mode. The structure of dihydrofolate reductase minimized in the presence of each of the other binding modes is very similar (not shown).

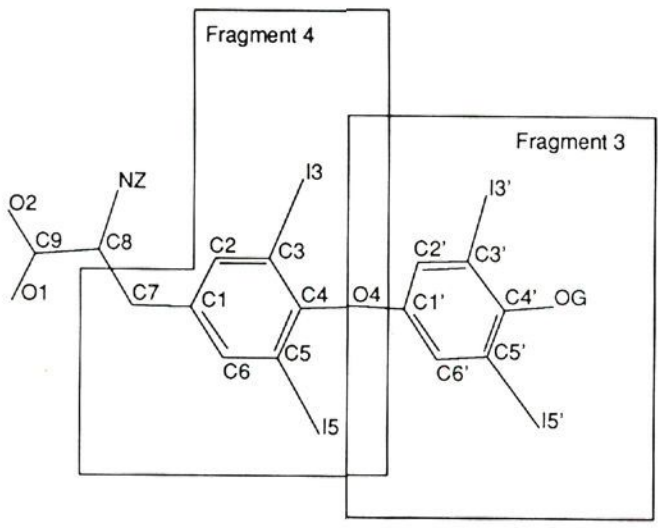

Figure 3. The thyroxine molecule showing division into fragment 3 and fragment 4.

Table III. Binding Modes for Prealbumin/Thyroxine

	fragment 3	fragment 4
atoms	10	10
initial matches[a]	503	480
orientations with no overlaps[b]	198	176
nonredundant orientations[c]	142	134
joined solutions (modes)[d]	19	
families[e]	6	

[a] We required a minimum of five atom–sphere pairs for a match with the distance error parameter C = 1.5 Å. [b] An overlap is defined as having any fragment atom within 2.5 Å of any receptor atom. [c] An orientation is considered redundant if it is within 0.1-Å rms deviation of any other orientation. [d] A joined solution was saved if the following distance conditions were met: O4 fragment 3 0.0–1.25 Å from O4 fragment 4; C1′ fragment 3 2.0–4.0 Å from C4 fragment 4. All atoms not specified above in either fragment must be more than 2.5 Å from all atoms in the other fragment. [e] Modes are in the same family if they differ by less than 1.0 Å, rms.

PAB tetramer in the absence of ligand has been determined at 1.8-Å resolution.[20] We used the Brookhaven Protein Data Bank set 2PAB.[18] The interaction of T4 with PAB has been examined crystallographically at low resolution[21] and been examined through model-building with molecular graphics.[4]

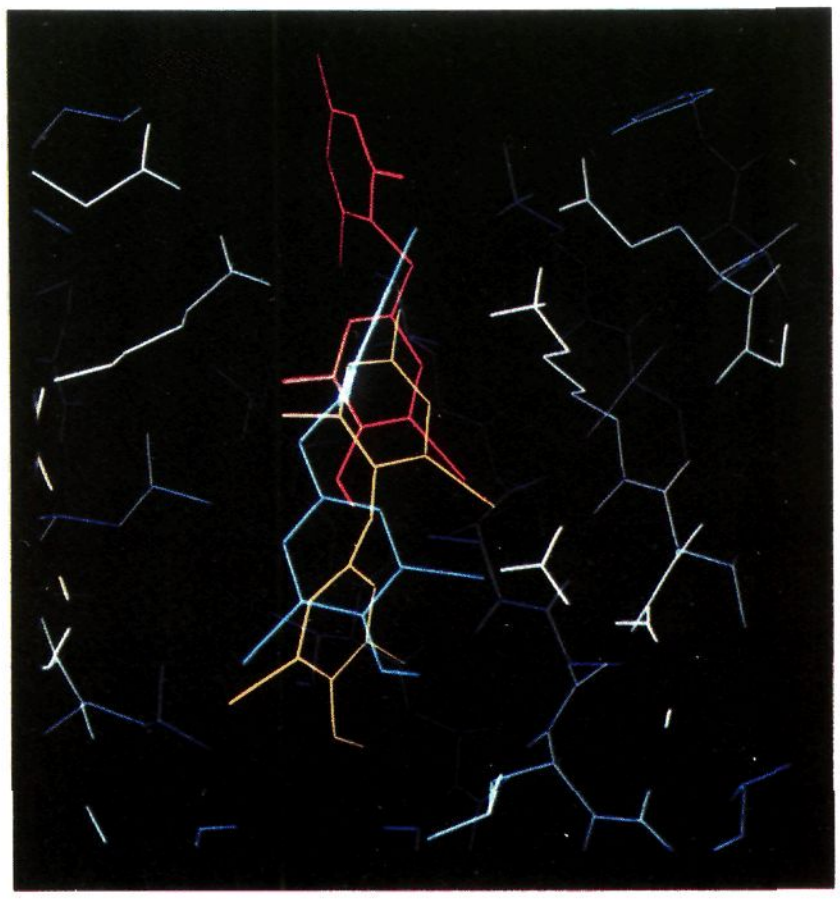

Figure 4. Selected thyroxine binding modes after energy minimization: model-built structure from Blaney et al.[4] (blue), family 4 (red), and family 6 (yellow). The structure of prealbumin is that minimized in the presence of thyroxine bound in the model-built mode. The structure of prealbumin minimized in the presence of each of the other binding modes is very similar (not shown).

PAB presents a more difficult case than DHFR for characterizing the shape of the binding site. The bottom of the prealbumin site is a deep and entirely reentrant surface. It was therefore impossible to characterize the site completely with use of only the contact surface. Kuntz et al.[9] did not encounter this problem because they included a crystallographically observed tightly bound water

(20) Blake, C. C. F.; Geisow, M. J.; Oatley, S. J.; Rerat, B.; Rerat, C. *J. Mol. Biol.* **1978**, *121*, 339.

(21) Blake, C. C. F.; Oatley, S. J. *Nature (London)* **1977**, *268*, 115.

Table IV. Energies of Binding Modes for Prealbumin/Thyroxine[a]

	rel energy, kcal/mol		rel energy, kcal/mol
X-ray crystal structure	0.0	family 4	7.0
family 1	2.0	family 5	12.0
family 2	3.0	family 6	−2.0
family 3	4.0		

[a] Energies were calculated with AMBER[15] using the new force field parameters.[16] Only those atoms accounted for by fragments 1 and 2 were included for the ligand and only those residues in prealbumin within 8 Å of the modeled position of thyroxine[4] were included in the calculation. The absolute energy of the modeled ligand after minimization was 659.2 kcal/mol.

molecule at the bottom of the site that provided contact surface. Since it is not clear a priori whether the water will be displaced by T4, we chose not to include the water in our analysis. In order to generate a complete representation of the binding site, we used both contact and reentrant points. Using more surface points to obtain a more complete and/or detailed representation of binding site shape is an option that is always open to the user without any change in the algorithm. Of course, this increases the computational time in the sphere-generation step and in the match step, since more spheres are produced.

The T4 molecule consists of two substituted phenyl rings attached by an ether linkage. There is a flexible amino acid substituent on one of the rings. The structure is illustrated in Figure 3. For the purposes of this calculation, the amino acid was omitted. The remaining portion is divided into two fragments called fragment 3 and fragment 4. The portion of T4 consisting of fragments 3 and 4 will be referred to as the ligand.

The results of matching are shown in Table III. Regenerating the ligand is done by requiring that the atom O4 in fragment 3 and O4 in fragment 4 be within 1.25 Å and that the distance between C1′ of fragment 3 and C4 of fragment 4 be 2–4 Å so that a reasonable C4–O4–C1′ angle be attained. We find 19 modes that can be divided into six families.

Again, representatives from each family were energy minimized with AMBER.[15] Those residues within 8 Å of the crystallographically determined position of T4 were included in the calculation. As before, position constraints at the termini were used, and explicit hydrogens were added. As with DHFR, prealbumin minimized to the same conformation except for some slight shifts of the side chains. In the case of PAB, a pair of symmetrically related lysine residues, which partly block the binding site shifted the most.

The energies are listed in Table IV. The comparative energies do not divide the families neatly into classes as in the DHFR case. There is one family that attains a lower energy than the X-ray mode. In that family, T4 is in an orientation similar to the X-ray mode, but displaced more deeply into the binding site. The OG atom of T4 forms hydrogen bonds with the protein; this accounts for its lower energy. In the crystal structure of the PAB–T4 complex, the critical hydrogen bonds are formed between the protein and the tightly bound water instead of with T4. Other families include T4 orientations with the OG pointing out toward the solvent, and modes where T4 is displaced further out of the binding site. Figure 4 shows two of the families and the modeled structure of Blaney et al.[4]

Discussion

In the cases presented here, the systematic search took into account all the degrees of freedom for docking rigid fragments. In each case, a ligand-binding geometry was found that was much like that observed crystallographically. Some other geometries, which were similar in energy, were found as well. These modes might not be found by manually docking the ligand.

While these results are encouraging, it is important to note the assumptions that limit the use of this method. The first assumption is that binding is determined primarily by shape complementarity. For example, if a ligand is held in place tightly by hydrogen bonds, but only loosely by steric forces, it is very unlikely that our method would find that binding geometry. A related assumption is the receptor site should have only small changes in shape when the ligand is bound. While the energy-minimization step allows a small amount of induced fit, the starting geometries for optimization are determined entirely by the initial atomic coordinates of the receptor. In our test cases, the coordinates represent the receptor in the absence of any ligand (PAB) or in the presence of a ligand bound in one particular geometry (DHFR).

The major advance in this paper is that the method is no longer limited to completely rigid ligands. We can treat those ligands that can be approximated by a small set of fragments connected by flexible linkages. We also note that it was not necessary for the shape representation of the ligand to take all atoms into account. In the cases presented here, a small flexible "tail" of each ligand was eliminated altogether, but useful results were still obtained. This is not surprising since in both cases the portion excluded from the ligand-shape representation does not contribute greatly to the overall shape complementarity of receptor and ligand. When this is the case, a portion of the ligand may be ignored in the matching process. The presence of such a piece may be taken into account later to exclude orientations that would place it in van der Waals contact with the receptor. As an example, orientations 3 and 4 in the DHFR test case could be excluded since they would leave no room for the glutamate portion of the intact methotrexate.

Even with these limitations, docking by means of shape complementarity has many potential uses. We are not limited to docking known ligands. A library of simple molecular shapes can be scanned to find which shapes most closely fit various subsites in a known receptor. Novel ligands can be built up by linking together the shapes in chemically reasonable ways. One need not have an X-ray structure of the receptor to characterize the shape of the receptor cavity for this application; the shape of the binding cavity can in some cases be approximated by comparing the shape of several active ligands. We are not even limited to docking flexible ligands into receptor cavities. One can imagine docking a flexible ligand into the shape of a rigid ligand. This could aid in finding common shape "pharmacophores" in seemingly unrelated molecules. We are currently exploring some of these applications.

Acknowledgment. We thank Dr. J. Blaney for providing the coordinates of the prealbumin/thyroxine model. We thank the Computer Graphics Laboratory (supported by NIH, Grant RR01081) for use of their facilities for parts of this project. Partial support was received from the National Institutes of Health (I.D.K.).