



A review of protein-small molecule docking methods

R. D. Taylor^{1,*}, P. J. Jewsbury² & J. W. Essex^{1,†}

¹Department of Chemistry, University of Southampton, Highfield, Southampton, SO171BJ, UK; ²AstraZeneca, Mereside, Alderley Park, Macclesfield, Cheshire, SK10 4TG, UK

Received 25 January 2002; accepted 18 April 2002

Key words: lead optimisation, molecular docking, protein-ligand complexes, scoring, virtual screening

Summary

The binding of small molecule ligands to large protein targets is central to numerous biological processes. The accurate prediction of the binding modes between the ligand and protein, (the docking problem) is of fundamental importance in modern structure-based drug design. An overview of current docking techniques is presented with a description of applications including single docking experiments and the virtual screening of databases.

Introduction

The number of algorithms available to assess and rationalise ligand protein interactions is large and ever increasing. Many algorithms share common methodologies with novel extensions, and the diversity in both their complexity and computational speed provides a plethora of techniques to tackle modern structure-based drug design problems [1]. Assuming the receptor structure is available, a primary challenge in lead discovery and optimisation is to predict both ligand orientation and binding affinity; the former is often referred to as ‘molecular docking’ [2]. The algorithms that address this problem have received much attention [3], indicating the importance of docking to a drug-design project. Owing to the increase in computer power and algorithm performance, it is now possible to dock thousands of ligands in a timeline which is useful to the pharmaceutical industry [4].

Despite the large size of this field, we have attempted to summarise and classify the most important docking methods. The principal techniques currently available are: molecular dynamics, Monte Carlo methods, genetic algorithms, fragment-based methods, point complementarity methods, distance geometry

methods, tabu searches and systematic searches. Algorithm examples and the test cases used to validate the models will be discussed for each of these approaches.

Large scale docking and virtual screening

Molecular docking is often used in virtual screening methods [5], whereby large virtual libraries of compounds are reduced in size to a manageable subset, which, if successful, includes molecules with high binding affinities to a target receptor. The potential for a docking algorithm to be used as a virtual screening tool is based on both speed and accuracy. This review will therefore highlight those docking methods that have been used in virtual screening applications.

Docking and de novo design methods

For the purpose of this review, a broad distinction is made between docking algorithms and *de novo* design methods. This is arguably subjective and in many cases significant overlap in methodology occurs between the two strategies. Examples of *de novo* design tools are BUILDER [6], CONCEPTS [7], CONCERTS [8], DLD/MCSS [9], Genstar [10], Group-Build [11], Grow [12], HOOK [13], Legend [14], LUDI [15], MCDNLG [16], SMOG [17] and SPROUT [18]. LUDI is given as an example of a *de novo* design tool applied to the docking problem.

*Present address: Astex Technology Ltd. 250 Cambridge Science Park, Cambridge, CB4 0WE, UK

†Author to whom correspondence should be addressed.
(j.w.essex@soton.ac.uk)

Search algorithms

Docking protocols can be described as a combination of two components; a search strategy and a scoring function. The search algorithm should generate an optimum number of configurations that include the experimentally determined binding mode.

A rigorous search algorithm would exhaustively elucidate all possible binding modes between the ligand and receptor. All six degrees of translational and rotational freedom of the ligand would be explored along with the internal conformational degrees of freedom of both the ligand and protein. However, this is impractical due to the size of the search space. For a simple system [19] comprising a ligand with four rotatable bonds and six rigid-body alignment parameters, the search space has been estimated as follows. The alignment parameters are used to position the ligand relative to the protein in a cubic active site measuring 10^3 \AA^3 . If the angles are considered in 10 degree increments and translational parameters on a 0.5 \AA grid there are approximately 4×10^8 rigid body degrees of freedom to sample, corresponding to 6×10^{14} configurations (including the four rotatable torsions) to be searched. This would require approximately 2 000 000 years of computational time at a rate of 10 configurations per second. As a consequence only a small amount of the total conformational space can be sampled, and so a balance must be reached between the computational expense and the amount of the search space examined.

The practical application of such an extensive search involves the sampling of many high energy unfavourable states which can restrict the success of an optimisation algorithm. In practice therefore, to sample such a large search space the computational expense is limited by applying constraints, restraints and approximations to reduce the dimensionality of the problem in an attempt to locate the global minimum as efficiently as possible. A common approximation in early docking algorithms was to treat both the ligand and target as rigid bodies and only the six degrees of translational and rotational freedom were explored. One of the first examples of such an algorithm is the early implementation of the program DOCK [20] (see Fragment-Based Methods). Although these methods have been successful in certain cases [21], there is a limitation to the rigid body docking paradigm in that the ligand conformation must be close to the experimentally observed conformation when bound to the target. Furthermore, numerous examples of conforma-

tional change of the target upon binding, for example the binding of cyclosporin A to cyclophilin [22], have led the drive to incorporate conformational flexibility into the search algorithm.

A common approach in modelling molecular flexibility is to consider only the conformational space of the ligand, assuming a rigid receptor throughout the docking protocol. The techniques used to incorporate conformational flexibility into a docking protocol will be discussed in some detail. However, the searching algorithm is only half the docking problem; the other factor to be incorporated into a docking protocol is the scoring function.

Scoring functions

Generating a broad range of binding modes is ineffective without a model to rank each conformation that is both accurate and efficient. The scoring function should be able to distinguish the experimental binding modes from all other modes explored through the searching algorithm. A rigorous scoring function will generally be computationally expensive and so often the function's complexity is reduced, with a consequential loss of accuracy.

Scoring methods can range from molecular mechanics force fields such as AMBER [23], OPLS [24] or CHARMM [25], through to empirical free energy scoring functions [26] or knowledge based functions [27].

The currently available docking methods utilise the scoring functions in one of two ways. The first approach uses the full scoring function to rank a protein-ligand conformation. The system is then modified by the search algorithm, and the same scoring function is again applied to rank the new structure.

The alternative method is to use a two stage scoring function. In this approach a reduced function is used to direct the search strategy and a more rigorous scoring function is then used to rank the resulting structures. These directed methods make assumptions about the energy hypersurface, often omitting computationally expensive terms such as electrostatics and considering only a few types of interaction such as hydrogen bonds. Such algorithms are therefore directed to areas of importance as determined by the reduced scoring function. Examples of directed methods are GOLD [28] and DOCK [29], and will be considered in more detail in the following sections.

A serious limitation in many existing scoring functions is the tendency to either neglect solvation effects

or use solvent models in a snap-shot fashion. A snap-shot method involves the generation of structures *in vacuo*, that are subsequently ranked with a scoring function that includes a solvent model. The search function is therefore directed to the conformational space which favours the *in vacuo* conformations. Furthermore, the structural role of bound solvent molecules and ions is often not considered, yet in the HIV-1 protease [30] system for example, it has been shown that explicit waters play an important role in ligand binding [31].

A brief description of the scoring and searching function will be given for each docking method in the following sections. The core components of the algorithm will be described, with a brief synopsis of the test cases used to validate the algorithms.

Molecular dynamics

There are many programs to perform molecular dynamics (MD) simulations such as AMBER [32] and CHARMM [25]. MD involves the calculation of solutions to Newton's equations of motions. Using standard MD to find the global minimum energy of a docked complex is difficult since traversing the rugged hypersurface of a biological system is problematic. Often an MD trajectory will become trapped in a local minimum and will not be able to step over high energy conformational barriers. Thus, the quality of the results from a standard MD simulation are extremely dependent on the starting conformation of the system.

This section focuses on novel MD techniques applied specifically to the docking problem to overcome the shortcomings of standard MD methodology.

Flexible ligands have been docked to flexible receptors in solution using MD simulations by Mangoni *et al.* [33], building upon the original work of Di Nola *et al.* [34]. The problems of obtaining adequate sampling are addressed by separating the centre of mass motion of the substrate from its internal and rotational motion. Separate thermal baths are then used for both types of substrate motion and receptor motion which permits local freezing of the various motion types.

Wang and Pak [35] have applied a new MD method to flexible ligand docking using a well jumping technique, where a scaling function is applied to the equations of motion to facilitate barrier crossing by effectively reducing the magnitude of the forces. Multicanonical molecular dynamics addresses the problem of limited conformational sampling and has been used

as a technique to dock flexible ligands by Nakajima *et al.* [36].

These methods operate on a single structure. However, it is common practice to generate a sub-ensemble of protein states, often using molecular dynamics, for use in docking studies. Such techniques have been summarised by Carlson and McCammon [37] where multiple protein structures are utilised rather than operating on a single flexible protein structure.

Monte Carlo methods

Monte Carlo (MC) methods are among the most established and widely used stochastic optimisation techniques. The combination of atomistic potential energy models with stochastic search techniques has produced some of the most powerful methods for both structure optimisation and prediction. A significant advantage of the MC technique compared with gradient based methods, such as MD, is that a simple energy function can be used which does not require derivative information. Furthermore, through a judicious choice of move type, energy barriers can simply be stepped over. The gradient based methods are often efficient at local optimisation, but have difficulty navigating a rugged hypersurface.

The standard MC method (more correctly, Metropolis MC [38]) involves applying random cartesian moves to the system and accepting or rejecting the move based on a Boltzmann probability.

Early implementations of AutoDock [39, 40] used Metropolis MC simulated annealing with a grid based evaluation of the energy, based on the AMBER force field, to dock flexible ligands into the binding pocket of a rigid receptor. The algorithm was originally tested on six complexes and was able to reproduce the experimental binding modes, although the lowest energy structures did not always correspond to the crystallographic conformation.

Prodock [41] uses a Monte Carlo minimisation technique to dock flexible ligands to a flexible binding site, using internal coordinates to represent the structures. This method differs from a standard MC procedure in that after each random move a local gradient-based minimisation is performed; the resulting structure is then accepted based on the Metropolis acceptance criteria. A grid based technique to evaluate the energy function is incorporated into the algorithm using Bezier splines [42], which produces a smooth function that can be differentiated; this property is

crucial to the local gradient-based minimisation. During the dock the magnitudes of the various potential energy terms are scaled to facilitate sampling. The independent scaling allows the selective reduction of barriers that restrict sampling. The size of each random move is determined from an assessment of the curvature of the hypersurface using the second derivative of the energy function. Thus large moves are attempted in areas of small curvature and small moves are attempted in areas of large curvature. Two force fields are implemented in Prodock, namely AMBER [23] and ECEPP/3 [43] along with a solvation model based on solvent exposed volume.

The MC method has been used to dock flexible ligands into a flexible binding site by Caflisch and co-workers [44]; this study built on previous work by Caflisch *et al.* [45] for docking an FKBP-Substrate complex. The first stage of the procedure places the ligand, at random, within the active site. This structure is then minimised *in vacuo* using a conjugate gradient minimiser with the CHARMM force field, allowing flexibility of the ligand and the protein. The Lennard-Jones and coulombic potentials are initially softened and gradually turned on throughout the course of the minimisation. This is repeated for 1000 seed structures. The seed structures are then ranked based on the potential energies calculated using the CHARMM force field. Solvation is included in the potential energy using a finite difference Poisson-Boltzmann (PB) term for the electrostatic contributions, calculated by UHBD [46], and nonpolar contributions are approximated by a weighted solvent-accessible area (SA) term. The MC method is then applied to the 20 structures with the lowest energy. This implementation of the MC method (referred to as Monte Carlo minimisation or MCM), is similar to the method adopted in the program Prodock [41]. MCM performs conjugate gradient minimisation after each random move. The minimised structures are then accepted based on the Boltzmann acceptance criteria. The energy for each MCM stage is again calculated using the CHARMM force field with the PB/SA solvent model. Each random move samples not only the position and orientation of the ligand but also a set of randomly selected dihedrals in the ligand and in the protein. This technique has been applied to three test systems and all three produced lowest energy structures within 1.4 Å RMSD of the crystallographic structures. Caflisch and co-workers also report the importance of allowing the protein to relax upon binding of the ligand, to discriminate near-native from non-native structures. This is

arguably one of the most ambitious docking projects to date.

Internal Coordinates Mechanics [47] (ICM) is a program to perform flexible protein-ligand docking and may be summarised as a MC minimisation method in internal coordinates. The algorithm initially makes a random move, which is one of three types; rigid body ligand move, torsion moves of the ligand, or torsion moves of the receptor side chain, using the biased probability methodology [48]. The side chain movement using this method is one of the defining features of the algorithm. The idea is to sample with a larger probability those regions of conformational space which are known *a priori*, based on previously defined rotamers [49], to be highly populated. This is achieved by making a normally distributed step in the vicinity of the low energy rotamer states for the protein side chains. Having made a random move, local minimisation of the ECEPP/3 [43] scoring function with a distance-dependent dielectric is performed using a conjugate gradient minimiser. An approximation for side chain entropy, loosely based around the statistical distributions of side chains, is then added to the minimised *in vacuo* ECEPP/3 energy. An electrostatic solvation term is then added to this energy, which is calculated using the MIMEL [48] approximation. This is a rapid approximation to the reaction field potential using the Born equation with a modification for many atoms. The modified ECEPP/3 energy is then used to test whether the structure is accepted or rejected, based on the Boltzmann criteria. A history mechanism has also been implemented to promote the discovery of new minima [50].

ICM has been applied to protein-ligand docking in the CASP-2 experiments [51]. For the 8 complexes tested only one produced an RMSD of 1.8 Å with respect to the crystal structure; the remaining test cases were only able to give, at best, an RMSD of 3 Å. However, in most cases the prediction was reasonable; on average 50% of the ligand was docked correctly.

MCDOCK [52] (version 1.0) applies a multiple stage strategy to dock a flexible ligand to a rigid receptor. The first stage of the docking places the ligand in the binding site. Random moves are then applied to the ligand to reduce the overlap of ligand and protein atoms. Metropolis MC with simulated annealing is then performed using a scoring function based on the CHARMM force field [25]. This is followed by a MC simulation which uses an adjustable temperature. In this method the temperature is increased if the acceptance ratio is too low, in an attempt to yield increased

sampling. MCDOCK was tested using 19 complexes, taken from the FlexX [53] optimisation test set. The RMSD between the binding modes predicted by MCDOCK and the experimental binding modes, for the non-hydrogen atoms of the ligand, ranged from 0.25 to 1.84 Å.

MC simulated annealing was applied to the docking problem using HIV-1 protease inhibitors by Bouzida *et al.* [54]. The AMBER force field was used with a desolvation correction based on the product of atomic charges and volume. To traverse efficiently the rugged energy hypersurface a soft-core smoothing function was used, for both the Lennard-Jones and electrostatic contributions to the potential energy. This methodology was used to dock two flexible ligands to a rigid X-ray structure of HIV-1 protease with some crystallographic waters retained as part of the rigid system. One of the docks reproduced the experimental binding mode. However, the second test case was not successful. The AMBER potential energy function was then exchanged for the piecewise linear potential function [55] (PLP). The PLP function is a simple model of ligand-protein interactions encompassing four terms: ligand and receptor non-bonded interaction terms (hydrogen bonds or steric clashes), internal torsion energies, and two penalty terms for leaving the active site and for internal clashes within the ligand. Using this scoring function the binding modes for both test cases were successfully reproduced.

Further MC simulations were performed using 10 different protein conformations for HIV-1 protease. The method consisted of randomly moving the ligand and calculating the score for this move, using the PLP function, between the ligand and all 10 protein complexes. The lowest energy from the 10 ligand-protein combinations was then used in the MC acceptance criteria to yield a frequency distribution of binding modes. This study attempted to rationalise the population of binding modes arising from the conformational changes in both the ligand and protein. They concluded, for two ligands, that there was a high correlation between protein conformation and predicted binding mode for one ligand but that the other case showed only a weak correlation.

DockVision [56] is another MC based docking method, using a rigid ligand and rigid receptor. The first stage of this docking algorithm generates a random ligand orientation. The MC method is then applied to the system, except the energy function is replaced by a geometric score for atomic overlap. This

is followed by an MC simulated annealing protocol using a simple potential energy function. The two stage docking procedure is then repeated for a large number of random ligand orientations. The ligand orientations generated by the MC dock are then clustered, based on a RMSD score. Two inhibitor complexes were used to test the protocol and in each case the binding geometry was correctly predicted. More recently this methodology has been applied to protein docking in CASP-2 experiments [57], achieving the second highest success rate.

QXP [58] performs MC flexible ligand/rigid protein docking, and is part of the FLO96 package. The Metropolis MC method is initially performed on the isolated ligand using only random dihedral moves (up to 360°). This is followed by rigid body rotations and translations to align the ligand onto guide atoms within the active site. These guide atoms are simply atoms in van der Waals contact with the binding site atoms. Having aligned the atoms within the active site, the MC method is applied to the ligand using only rigid body rotations and translations. Conjugate-gradient minimisation is then performed on the ligand torsions followed by Metropolis MC on the ligand torsions. In this method a grid representation of the receptor is used. The scoring function uses the AMBER force field with short non-bonded cut-offs and a distance dependent dielectric. The original test set consisted of 12 ligand-protein complexes, with a maximum of 24 rotatable ligand dihedrals. The ligand was flexible and the receptor rigid, with single important water molecules retained in three of the complexes. Their results were compared with energy minimised structures; 11 ligands gave an RMSD of less than 0.76 Å.

Affinity [59] is commercial program using Monte Carlo simulated annealing with a grid representation for the non-moving parts of the system [60] and an implicit representation of solvation effects [61]. Another commercial program is Glide [62] which uses a hierarchical filter to rapidly score hydrophobic and polar contacts, followed by Monte Carlo sampling with the ChemScore [26] scoring function.

Genetic algorithms and evolutionary programming

Since their inception, genetic algorithms (GA) have increased in popularity as an optimisation tool. It should be noted that GAs (and evolution programming (EP)) require the generation of an initial population

whereas conventional MC and MD require a single starting structure in their standard implementation.

The essence of a GA is the evolution of a population of possible solutions via genetic operators (mutations, crossovers and migrations) to a final population, optimising a predefined fitness function. Degrees of freedom are encoded into genes or binary strings and the collection of genes, or chromosome, is assigned a fitness based on a scoring function. The mutation operator randomly changes the value of a gene, crossover exchanges a set of genes from one parent chromosome to another, and migration moves individual genes from one sub-population to another.

GOLD [28] is a docking program that uses a GA search strategy and includes rotational flexibility for selected receptor hydrogens along with full ligand flexibility. Gene encoding is used to represent both rotatable dihedrals and ligand-receptor hydrogen bonds. A GA move operator is subsequently applied to parent chromosomes that are randomly chosen from the existing population with a bias towards the fittest members. The ligand-receptor hydrogen bonds are subsequently matched with a least squares fitting protocol to maximise the number of inter-molecular hydrogen bonds for each GA move. As a consequence the GA structure generation is biased towards inter-molecular hydrogen bonds. However each structure is ranked based on a more complex fitness function. The fitness (or scoring) function is the sum of a hydrogen bond term, a 4–8 inter-molecular dispersion potential and a 6–12 intra-molecular potential for the internal energy of the ligand. Each complex was run using an initial population of 500 individuals divided into five equal sub-populations, and migration of individual chromosomes between sub-populations was permitted. A single GA run used 100 000 genetic operations and 20 GA runs were performed. Finally, the solution with the highest fitness score was compared with the crystallographic binding mode.

The GOLD validation test set is one of the most comprehensive (comprising 100 different protein complexes) of all of the docking methods reviewed, and achieved a 71% success rate based primarily on a visual inspection of the docked structures. 66 of the complexes had an RMSD of 2.0 Å or less, while 71 had an RMSD of 3.0 Å or less. The omission of hydrophobic interactions and a solvent model may explain some of the docking failures which included highly flexible, hydrophobic ligands, and those complexes containing poorly resolved active sites. However, recent extensions to GOLD [63] include the

addition of hydrophobic fitting points that are used in the least squares fitting algorithm to generate the ligand orientation.

AutoDock 3.0 [64] uses a genetic algorithm as a global optimiser combined with energy minimisation as a local search method. In this implementation of AutoDock the ligand is flexible and the receptor is rigid and represented as a grid. The genetic algorithm uses two point crossover and mutation operators. For each new population a user determined fraction undergo a local search procedure using a random mutation operator where the step size is adjusted to give an appropriate acceptance ratio. The fitness function comprises five terms: a Lennard-Jones 12-6 dispersion/repulsion term; a directional 12-10 hydrogen bond term; a coulombic electrostatic potential; a term proportional to the number of sp^3 bonds in the ligand to represent unfavourable entropy of ligand binding due to the restriction of conformational degrees of freedom; and a desolvation term. This scoring function is based loosely around the AMBER force field from which protein and ligand parameters are taken. The desolvation term is an inter-molecular pairwise summation combining an empirical desolvation weight for ligand carbon atoms, and a pre-calculated volume term for the protein grid. Each of the five terms are weighted using an empirical scaling factor determined using linear regression analysis from a set of 30 protein-ligand complexes with known binding constants. The algorithm was originally tested on seven complexes, and for these test examples all lowest energy structures were within 1.14 Å RMSD of the crystal structure.

DIVALI [65] uses an AMBER-type potential energy function with a distance dependent dielectric and a genetic algorithm search function to dock four complexes. The receptor was modelled as a rigid entity and consequently a grid based energy evaluation of ligand-protein interactions was performed to assess the fitness function. An additional masking operator is used that fixes part of the population which is associated with translational space so that subpopulations search different regions of the active site. Three out of four of the test complexes gave an RMSD of 1.7 Å or less.

The program DARWIN [66] combines a GA and a local gradient minimisation strategy with the CHARMM-AA molecular mechanics force field, for flexible docking of three protein-carbohydrate complexes. Binary encoding is used to describe a starting ligand conformation and position; the potential energy is then locally minimised using a gradient

method and the chromosome fitness is scored using the CHARMM-AA potential energy function. The populations are then modified by standard mutation and crossover operators while the protein is held rigid. Solvent contributions are assessed using a modified version of the program DelPhi [67] to yield finite difference solutions to the Poisson-Boltzmann equation. Although the search algorithm was able to optimise the energy landscape, certain structures were obtained with energies lower than the experimental binding mode. The false positives produced were thus attributed to limitations in the scoring function. Including specific explicit waters in the binding site increased the success of the program. The authors further note the dynamic nature of the complexes, and that multiple binding modes is a reasonable reflection of reality and not an artefact of the force field.

Judson *et al.* [68] were one of the first to report the application of a GA to the docking problem. A flexible ligand was used with interacting sub-populations and a gradient minimisation during the search. The method was tested by docking Cbz-GlyP-Leu-Leu into thermolysin and produced conformations which were close to the experimental binding mode, although in some cases the energies were lower than the crystal conformation.

Gehlhaar *et al.* [55] have applied evolutionary algorithms to flexible ligand docking in an HIV-1 protease complex, using the previously described PLP scoring function [55]. An initial population is generated and the fitness of each member is evaluated based on the scoring function. The fitness of the members are then compared with a predetermined number of opponent members chosen at random. The members are then ranked by the number of wins, and the highest ranking solutions are chosen as a new population. All surviving solutions are used to produce offspring with a mutation operator, such that the population size is constant. This protocol is repeated until a user defined number of iterations is exceeded; a conjugate gradient optimisation is then performed on the best member. Interestingly, for this test case, previous docking attempts have failed. This failure was attributed to high energy barriers [69]. Consequently, the repulsive term was slowly turned on through the course of the simulation, in an analogous fashion to MC simulated annealing. 100 simulations were run and the crystal structure was reproduced 34 times with a maximum RMSD of 1.5 Å; these solutions were the lowest energy docks.

Fragment-based methods

The broad philosophy of fragment based docking methods can be described as dividing the ligand into separate portions or fragments, docking the fragments, followed by the linking of fragments. These methods require subjective decisions on the importance of the various functional groups in the ligand, which can result in the omission of possible solutions, due to assumptions made about the potential energy landscape. Furthermore, a judicious choice of base fragment is essential for these methods, and can significantly affect the quality of the results.

The docking of fragments and the subsequent joining of the docked fragments has been widely used in *de novo* design methods. Although, for this review, only a few *de novo* programs have been considered, there is a considerable overlap of methodologies.

One of the most popular programs to perform fragment docking is the incremental construction algorithm FlexX [53]. The initial phase is the selection of the base fragment for the ligand from which possible conformations are formed based on the MIMUMBA [70] torsion angle database. As with all fragment based methods the choice of base fragment is crucial to the algorithm; it must contain the predominant interactions with the receptor. Early implementations of FlexX required manual selection of this base fragment but this process has been subsequently automated [71]. Following the selection of the base fragment an alignment procedure is performed to optimise the number of favourable interactions. These interactions are based primarily on hydrogen bond geometric constraints but also include hydrophobic interactions. In this stage, the base fragment is considered rigid, and three sites on the fragment are mapped onto three sites of the receptor. All geometrically accessible receptor triangles are then clustered and the superposition of ligand triplets onto the receptor is performed using the method of Kabsch [72]. Overlaps are removed and energies are then calculated for the base fragments using Böhm's [73] function. Following this base fragment placement the ligand is built in an incremental fashion, where each new fragment is added in all possible positions and conformations. Intra-molecular and inter-molecular overlaps are then removed and the placements are ranked, from which the best solutions are subjected to a clustering protocol. The highest rank solution from each cluster is then used in the next iteration. This process is repeated until the complete

ligand is built, and the final structures are scored using the empirical scoring function.

The algorithm was originally validated on 19 complexes. Out of the 19 test cases, 10 of the docked complexes, with the best score, reproduced the experimental binding mode with a maximum RMSD of 1.04 Å. However, the experimental binding mode was reproduced for all 19 complexes, but these structures were not necessarily the lowest in energy. Recent extensions to FlexX include placing explicit waters into the binding site during the docking procedure [74] using pre-computed water positions. This was tested with 200 protein-ligand complexes and improvements were observed for some targets, such as HIV protease. Furthermore, receptor flexibility has been included using discrete alternative protein conformations [75].

The program DOCK (version 4.0 [29]) can be summarised as a search for geometrically allowed ligand-binding modes using several steps that include: describing the ligand and receptor cavity as sets of spheres, matching the sphere sets, orienting the ligand, and scoring the orientation. New extensions to the protocol involve combining the bipartite graphs consisting of protein and ligand interaction sites, into a single docking graph where each node now represents a pairing of an atom with a site point. Clique detection is then implemented based on the methodology of Bron and Kerbosch [76]. The technique has been previously evaluated [77, 78] and was found to be the most efficient methodology for finding cliques which encode maximal pairs of interactions between matching sites. Having generated multiple orientations an inter-molecular score is calculated based, on the AMBER force field [79] where receptor terms are calculated on a grid.

DOCK 4.0 includes ligand flexibility using a modified scoring function which incorporates an intramolecular score for the ligand [80]. In this version the docking is fragment based; a ligand anchor fragment is selected and placed in the receptor, followed by rigid body simplex minimisation. The conformations of the remaining parts of the ligand are searched by a limited backtrack method and minimised. This protocol was tested on 10 structures; 7 docked complexes reproduced the crystal structure with a maximum RMSD of 1.03 Å and the remaining 3 were within 1.88 Å.

Although DOCK is included as a fragment based method, this is only one of several modes of operation. An alternative mode of operation is the docking of multiple random ligand conformations [81, 82]. This method generates a user-defined number of conform-

ers as a multiple of the number of rotatable bonds in the ligand. If the total number of user-defined conformers is greater than the number of conformations possible, based on a set of dihedral rules, then a systematic search is performed. Otherwise the required number of conformers are generated by assigning random dihedral values. These conformers are then docked independently. The search algorithms available in DOCK 4.0 have recently been reviewed by Ewing *et al.* [82]. Further extensions to DOCK have included incorporating protein flexibility using ensembles of protein structures [83] and the inclusion of a GB/SA [84] continuum model into the scoring function [85].

The program LUDI [15] fits ligands into the active site of a receptor by matching complementary polar and hydrophobic groups. The program operates a three stage protocol by first calculating protein and ligand interaction sites that are discrete positions in space, which can either form hydrogen bonds or fill hydrophobic pockets. These sites are derived in one of three ways: from non-bonded contact distributions based on a search through the Cambridge Structural database, from a set of geometric rules, or from the output from the program GRID [86]. GRID calculates binding energies for a given probe, e.g. carbonyl, amide, aromatic carbon, with a receptor molecule. The next stage involves the fitting of molecular fragments onto the interaction sites. If the distance between interaction sites on the receptor is within a user defined range of those on the ligand then the algorithm attempts a fragment fit using an RMSD superposition algorithm [72], which fits complementary ligand groups onto the receptor. A fragment position is then accepted if the RMSD is within a specified range. Limited internal flexibility of the ligand can be incorporated using several conformers of a fragment in a library. The final stage of the protocol involves the joining of fragments using a database of bridge fragments and the previous fitting algorithm. LUDI was originally validated by predicting modifications to HIV protease inhibitors and DHFR inhibitors [87]. Two fragments were predicted as modifications to an inhibitor of HIV protease, both of which were found experimentally to yield substantially improved binding. Seven structures were obtained as modifications to the anti-cancer drug methotrexate complexed with dihydrofolate reductase (DHFR). Experimental data showed an improvement in binding for one of these seven complexes predicted by LUDI.

ADAM [88] is another fragment building algorithm for flexible docking which can be summarised as alignment of a fragment based on hydrogen bonding motifs, followed by energy minimisation in the AMBER program. The initial stage requires the user to choose a ligand fragment. Hydrogen bonding dummy atoms are then calculated based on distance and direction from potential hydrogen bond groups in the protein. These dummy atoms are matched with all potential hydrogen bond atoms in the ligand fragment, for a predetermined set of fragment conformations generated systematically by rotating the torsion angles. Only low energy conformations for the ligand are used, as determined by the intra-molecular energy calculated using a modified version of the AMBER force field. The matching of the ligand fragment to the protein hydrogen bonding sites is performed using an algorithm based on the superposition method by Kabsch [72]. Minimisation of the hydrogen bonding fragment is then performed using the Simplex [89] algorithm, again using the modified AMBER potential energy. The conformations of the rest of the ligand are then scanned first in large and then small angle intervals, to determine the low energy conformations. The final ligand conformations are then minimised using AMBER. The methodology was applied to two protein systems: dihydrofolate reductase (DHFR) (with three ligands) and ribonuclease T₁ (one ligand). In both cases, an RMSD of less than 1 Å was obtained for the lowest energy structures.

Hammerhead [19] is another fragment based docking program which is similar to ADAM. Head fragments are initially generated by dividing the ligand into sections. A systematic search is then used to generate a diverse set of fragment conformations. These sets are then aligned to hydrophobic or hydrogen bonding groups in the protein, and scored with an empirical measure of binding affinity. At each stage of alignment a steepest-descent optimisation is performed to enhance the score. The highest scoring fragments are considered head fragments. To each head fragment, the tails (or remaining fragments) are added one at a time. This is achieved by aligning the tail with the protein, and then merging the tail with the head fragment. The best scoring orientations are then retained for the addition of the next fragment. It should be noted that FlexX, ADAM and Hammerhead all use a scoring function to rank intermediate stages of the construction algorithm, to control a possible combinatorial explosion. The algorithm was tested on four complexes and produced an RMSD less than 1.7 Å

in each case. The method was also tested on a set of 80 000 compounds binding to streptavidin; biotin (a ligand with a high affinity binding to streptavidin [90]) was predicted as the highest-scoring ligand.

The virtual screening algorithm SLIDE [91] and its precursor SPECITOPe [92] also use a fragment approach for molecular docking. Anchor fragments are chosen which contain triplets of matching points, where a matching point is a hydrogen bond donor, acceptor or hydrophobic ring centre. The best matches between triplets of matching points on the ligand with matching points on the protein are then calculated based on chemistry and geometry. A least squares superposition algorithm is then used to place the fragment into the binding site. The rest of the ligand is added to the base fragment using the original database conformation. Any intermolecular overlaps are subsequently removed by rotating protein side-chains and ligand dihedrals based on mean-field theory [93]. The complexes are then scored based on hydrogen bonds and hydrophobic complementarity. This simple score was parametrised using the ChemScore training set of protein-ligand complexes. SLIDE has recently been applied to database screening for HIV-1 protease ligands [94] which included an explicit water molecule in the binding pocket. Using a database containing 15 known ligands SLIDE was able to identify approximately half of the true hits.

Point complementarity methods

Docking ligands to the binding site of a receptor is often performed using points of complementarity between the protein and ligand. Arguably, many of the fragment based docking algorithms could also be included in this category, although a distinction has been made between algorithms that treat the ligand as a complete entity throughout the docking method, and those where the ligand is divided into fragments.

Jiang and Kim [95] proposed the 'Soft docking' method which implicitly allows for small conformational changes in the ligand and receptor. The molecular surface is represented as a series of surface cubes (along with the surface normal) and volume cubes i.e. cubes inside the molecule. The cube description of the ligand is then rotated and translated to obtain the maximum number of matches between ligand and protein surface cubes, minus the number of volume cube overlaps. A surface cube match must satisfy not only an overlap between cubes in the two molecules

but also the surface normals must be approximately in opposite directions. By allowing penetration of surface cubes the molecular surface is effectively softened. The docking solutions are then clustered based on translation vectors and rotation angles. The average value for each cluster is then scored using a geometric sum of atom descriptors, that are based on charges, hydrogen bond donors/acceptors and hydrophobicity. The method was tested on four complexes using the bound conformations and in all cases good agreement with experiment was seen for the solutions with the most favourable score. A further two examples were tested using unbound conformations. The first example used the free form of both the ligand and the protein; the best docking solution was ranked sixth. The RMSD for this solution with the X-ray structure of the complex was 2.56 Å, for the ligand. This was stated to be reasonable by the authors, since the RMSD from a least-square fitting of the free ligand to the complexed ligand was 1.80 Å. In the second example a protein-protein dock was attempted. The complexed and uncomplexed form of two proteins were docked and the solution which was ranked eighth achieved an RMSD of 1.55 Å.

FTDOCK [96] is a rigid docking protocol based on shape complementarity with a model for the electrostatic interactions. A grid is placed over the protein and over the ligand and each grid point is labelled either open space or inside the molecule (i.e. the value 0 or 1). The protein is further divided into grid points on the surface of the molecule (value of 1) or buried within the molecule (value of -15). The function to be optimised, by rigid body translation and rotation of the ligand, is referred to as the correlation function. This function is the product of a grid point in the ligand with the corresponding grid point of the protein, summed over all points. A high correlation score denotes good surface complementarity between the molecules. This function is rapidly calculated using fast fourier transforms. An electrostatic score is also used based on a rudimentary atomic charge for the protein and ligand. This additional score is used as a binary filter to remove false positives with a high positive electrostatic contribution. The method was tested on six enzyme/inhibitor and four antibody/antigen complexes, using the apo-protein structure rather than the complexed forms. In all but one of the test cases, an RMSD of ≤ 2 Å was achieved for the C $^{\alpha}$ of the protein and ligand. However these were seldomly the highest ranking solutions. This method has also been tested in the CASP-2 protein-protein docking competition [97].

Jackson *et al.* [98] have refined the protein-protein interfaces for structures generated by FTDOCK, using the program MULTIDOCK. In this work the protein was modelled using a multiple copy representation of side chains based on rotamer libraries, that were subsequently refined by a mean field method followed by energy minimisation.

Both FTDOCK and the 'Soft docking' method were originally developed for protein-protein docking. Owing to the computational expense of this problem it is necessary to use the rigid body approximation. However, this approximation is a limitation in ligand-protein docking.

LIGIN [99] uses a surface complementarity approach to dock rigid ligands into a rigid receptor. From a set of random starting points a complementarity function is maximised. This function is a sum of atomic surface contacts that are weighted according to whether or not the interactions are favourable. The ligand structures generated from the first stage are then moved to maximise hydrogen bond distances between the ligand and the protein. This method was applied to the CASP-2 ligand docking test set [100]. Seven docking predictions were made; two ligands were docked to within 1.8 Å of the crystallographic structure and three other test cases could only reproduce the buried parts of the ligand.

SANDOCK [101] addresses the protein-ligand docking problem using both shape and chemical complementarity. The ligand is fitted onto the accessible surface of the protein binding site by a distance matching algorithm. If a certain number of interaction distances can be found for the ligand that, within a tolerance, match the surface cavity for the protein binding site, a transformation of the ligand into the protein binding site is attempted. The tolerance for the distance complementarity is adjusted to favour hydrophobic and hydrogen bonding interactions between the ligand and protein. The ligand orientations are then scored based on a geometric function which penalises close contact of atoms, coupled with a hydrophobic and hydrogen-bonding score. The method was tested on a thrombin-ligand complex producing an RMSD of 0.7 Å.

FLOG [102] (Flexible Ligands Oriented on a Grid), selects ligands, from a database, complementary to a receptor of known three-dimensional structure. The algorithm philosophy is similar to DOCK [103], but uses up to 25 explicit conformations of the ligand to allow for some flexibility. Distances between favourable sites of interaction in the protein

are calculated along with atom-atom distances of the ligand. A clique-finding algorithm is then used to calculate the sets of distances, that can occur simultaneously, between the ligand and the protein interaction sites. The ligand is then superimposed onto the protein interaction sites and optimised using a simplex algorithm. Each orientation is then scored (using a grid representation for the receptor) with a function that includes explicit terms for van der Waals, electrostatics, hydrogen bonding and hydrophobic interactions. This methodology has been tested on dihydrofolate reductase to select known inhibitors from a database. The highest scoring conformation from FLOG did not match the crystallographic binding mode. However, the algorithm was designed to find all potentially interesting compounds in a database and was shown to successfully identify known inhibitors from a database of drug like compounds.

Distance geometry methods

Distance geometry methods like those developed by Leach *et al.* [104] determine the binding modes between protein and ligand considering only hydrogen bonding. This method samples conformational space identifying plausible binding modes which are then used to direct an embedding algorithm.

A more recent application of distance geometry to flexible ligand docking to a rigid receptor is the program DockIt [105]. The active site is represented as a set of spheres and ligand conformations are generated within this site using distance geometry. The resulting structures are then scored using either the PMF [27] or PLP [55] scoring functions.

Tabu searches

PRO_LEADS [106, 107] is a docking algorithm which may be described as a stochastic evolution of the system using a tabu search with a generalised scoring function, ChemScore. An initial random ligand conformation is generated (referred to as the *current solution*) and then scored. Random moves are then applied to the ligand to generate a population of solutions (typically 100 solutions). These solutions are then scored and ranked in ascending order. The highest ranking solution is then accepted as the new *current solution*, (assuming it is the lowest energy so far). A new random population is then generated from this

new *current solution*, and the process is repeated for a user defined number of iterations. However, to ensure diversity of the *current solution* a tabu list is used. This list stores the orientations of the last 25 *current solutions*. If the lowest energy solution from the ranked population, is the lowest energy so far, it is always accepted as the new *current solution*. However, if it is not the lowest energy so far, the best non-tabu solution is then used. A move is considered 'tabu' if it is within an RMSD of 0.75 Å of any of the solutions stored within the tabu list. This tabu list is updated on a first-in first-out basis. The scoring function uses simple contact terms to estimate lipophilic and metal-ligand binding contributions, with an explicit term for hydrogen bonds and a term which penalises flexibility. The lipophilic and metal terms are based on atom-atom contacts, while the hydrogen bond term is a function of both distance and angle. The flexibility term is derived from the rotatable bond count, modified by the chemical environment of the bond. The function has been parametrised on an 82 complex training set with known binding affinities.

Flexible ligand docking was originally performed using PRO_LEADS with a validation test set of 50 ligand-protein complexes and achieved a success rate of 86% of the solutions within an RMSD of 1.5 Å of the crystallographic structures. This is also one of the largest validation test sets for a docking algorithm. The program has also been applied to virtual database screening [108]. For this application, the method was first tested on 70 ligand-receptor complexes and 79% of the solutions were within 2.0 Å. A database of 10000 randomly chosen druglike molecules were docked into three target receptor structures. For each receptor, known ligands were included in the database. Enrichment factors of approximately 8 were achieved for the top 10% of the docking solutions for all three receptors.

Systematic searches

EUDOC [109] is a systematic search of rigid body rotations and translations of a rigid ligand within a rigid active site. It is based on the earlier program SYSDOC [110] which uses the fast affine transformation to perform the systematic search. The intermolecular energy is calculated by the AMBER-AA force field with a distance dependent dielectric. The method was applied to virtual screening of farnesyl-transferase inhibitors from which 21 hits were found.

Four inhibitors from the 21 hits were deemed to be good inhibitors; no true hits were identified from 21 randomly selected compounds.

Multiple method algorithms

Combining different docking techniques into a single strategy is a useful method to increase the effectiveness of a docking protocol. Often a two stage strategy is adopted using an initial computationally inexpensive method, followed by a time consuming yet more accurate method to generate a final docking solution.

We have developed a novel flexible docking algorithm [111] which samples extensively the protein and ligand conformations, and uses an 'on-the-fly' continuum solvent model. The method may be summarised as a hybrid approach where the first stage is a rapid dock of the ligand to the protein binding site, based on deriving sets of simultaneously satisfied hydrogen bonds using graph theory and a recursive distance geometry algorithm [104]. The output structures are then reduced in number by cluster analysis based on distance similarities. These structures are then submitted to a modified Monte Carlo algorithm [112] which uses the AMBER-AA molecular mechanics force field with the Generalised Born/Surface Area (GB/SA) continuum model [84]. This solvent model is not only less expensive than an explicit representation but also yields increased sampling. Sampling was also increased using a rotamer library to direct some of the protein side chain movements along with large dihedral moves. Furthermore, a soft-core function for the non-bonded force field terms was used, enabling the potential energy function to be slowly turned on throughout the course of the simulation.

The docking procedure was optimised on a single complex and validated on a further 14 complexes. For each complex the flexible ligand was docked into the receptor both with and without protein side chain flexibility. In the rigid protein dock 13 out of the 15 test cases were able to find the experimental binding mode; this number was reduced to 11 in the flexible protein dock. In these instances, although the experimental binding mode could not always be uniquely identified, in the majority of cases this mode was present in a cluster of low energy structures that were energetically indistinguishable.

Leach [113] proposed an algorithm to explore both conformational flexibility of the ligand with the conformational degrees of freedom of the amino acid

side-chains. In this work the flexible docking problem was restricted to finding the lowest energy combination of amino acid conformations with pre-computed low energy ligand conformations, for discrete ligand orientations. Side-chain conformations were assigned from a library of minimum energy structures based on discrete amino acid rotamer states [114]. The Dead End Elimination algorithm was used to remove unfeasibly high energy rotamer states from the library of rotamer states. Then the tree search algorithm A* was used to find the minimum energy combinations for rotamer states of the protein with the ligand. Each solution was scored using the AMBER force field. The method was applied to two test complexes. Only one test case was successful in locating the lowest energy conformation of both the ligand and protein side chains. An interesting observation was the larger number of rotamer combinations found for the complexed protein compared with the isolated structure, suggesting that more conformational states are accessible in the presence of a ligand.

MC simulated annealing with the Dead End Elimination algorithm for side chain optimisation [115] has been used to predict structural effects in HIV-1 Protease. In the initial stage 25–50 independent MC simulated annealing simulations are performed to dock a flexible ligand into a static representation of the protein. For each solution, the protein side chains are optimised using the Dead End Elimination [114] procedure. The energy is then minimised for all of the docking solutions using a conjugate gradient minimiser. The final solutions are then ranked using the AMBER [116] force field with the GB/SA [84] solvation model. Two docked ligands achieved an RMSD of less than 1.04 Å with the bound crystal structure, with 70% of the residues correctly predicted, starting from the wild-type HIV-1 protease structure. The protocol was then used to predict effects for the HIV-1 protease mutants. The ligand conformations were identified to within an RMSD of 1 Å, with 80% of the residue conformations correctly predicted. The authors also came to the conclusion that protein side chains are insensitive to the orientation of the ligand.

A recent technique proposed by Wang *et al.* [117] follows a multi-step strategy for flexible ligand docking which separates conformational and orientational space. Low energy conformers of the unbound ligand are selected using a systematic search and rigid docking is then applied using the program DOCK [20]. The low energy orientations are then minimised using first a steepest-descent minimisation, followed by conju-

gate gradient minimisation. In both cases the AMBER [23] force field was used to score the structures, and the receptor was held rigid. The torsion angles of the ligand are then refined, in the binding site, using a systematic search. The final stage of the method is a short period of MD based simulated annealing. The method was applied to 12 ligand-protein complexes; the lowest energy structures for each complex were within an RMSD of 2.01 Å with the crystal structure.

A tabu search followed by structure refinement with a MC simulated annealing protocol was implemented by Price and Jorgensen [118]. The method was used to determine starting conformations for MC/free energy perturbation calculations. The tabu procedure was reported to find the crystal structure with greater frequency and at a lower computational cost compared with a simulated annealing protocol. This observation was confirmed by Westhead *et al.* [119].

A two stage strategy was also used by Hoffmann *et al.* [120] which can be summarised as rapid docking using FlexX [53], followed by minimisation and ranking of the structures using the CHARMM program [25]. 10 different complexes were tested, and the two stage strategy significantly improved the results, compared with the results obtained from only the fast docking stage.

Mining Minima optimisation was applied to the docking problem by Gilson *et al.* [121]. An initial random structure was generated and random moves were applied using a focussing method. This ensures new solutions are generated around the best scored members from previous solutions. GA based crossover moves were also used, along with moves based around the poling [122] method and tabu search. This ensures that no two solutions have both the same internal and rigid body coordinates. The function to be optimised was the CHARMM force field with a distance-dependent dielectric. Using 5 complexes selected from the PRO_LEADS [106, 107] test cases the algorithm was able to successfully dock all ligands. A further 14 ligands were also docked successfully. Using an additional 9 complexes from the GOLD test set, (all unsuccessful docks using GOLD) the Mining Minima method was able to successfully dock 6.

Comparison of techniques

There have been a number of attempts to compare different docking protocols, where the comparison is often divided into an assessment of the search al-

gorithm and the scoring function. Each attempt at an unbiased comparison has clearly stated that it is very difficult to rank and assess different optimisation strategies, particularly with increased algorithm complexity. The authors also note the performance of any docking protocol will not only depend on the core algorithms but also the time invested in parametrisation and optimisation of the methodology.

A recent analysis by Vieth *et al.* [123] of prevalent searching techniques using the CHARMM force field compared MD, MC, and a GA in the docking of 5 complexes. Their analysis showed that the performance of the different techniques depended on the size of the binding site. They suggested that MD provided the best efficiency in docking structures in a large space, whereas the GA is best for a small search space. This conclusion seems to contradict other results; for example, AutoDock3.5 [64] has recently transferred from MC/simulated annealing to a GA searching function, based on the ability of a GA to efficiently sample large search spaces. Furthermore, Vieth *et al.* used an annealing protocol in both the MC and MD applications, whereas the GA equivalent utilised a high mutation rate in the initial stage followed by a high crossover rate in the second stage. The latter corresponded to a lower temperature anneal in the MC and MD analysis. The MC algorithm was also shown to perform reasonably well on both the large and small search spaces. In an accompanying article energy functions have been assessed for flexible docking [124] by modifying the CHARMM [25] parameter set with an MD simulated annealing protocol on 5 ligand-protein complexes. The largest improvements in the docking efficiency came from the reduction of the van der Waals repulsion and a reduction of surface charges.

Another analysis compared four search algorithms, namely MC/simulated annealing, GA, EP and tabu search applied to flexible ligand docking [119] for 5 complexes, using the PLP scoring function developed by Gehlhaar *et al.* [55]. The GA gave the lowest median energies, but the tabu located the global minimum more reliably. The tabu search was therefore the searching algorithm employed by the authors in the program PRO_LEADS [106, 107].

Foreman *et al.* [125] compared simulated annealing, MC and the convex global under-estimator (CGU) as optimisation search methods for finding the global minimum on an energy hypersurface. CGU is similar to the method applied in Mining Minima [121], and involves choosing random points on the energy

hypersurface, around which the landscape is approximated as a parabola. The CGU method was shown to reach lower energies for a protein folding energy landscape compared with a simulated annealing/MC protocol and the standard MC method.

Dock, FlexX and GOLD have been evaluated in combination with seven different scoring functions by Bissantz *et al.* [126]. Two protein targets were selected (thymidine kinase (TK) and estrogen receptor) along with two random databases of 990 ligands with a further 10 true hits in each. GOLD gave the best RMSD solutions and best ranking of ligands for both TK and the estrogen receptor out of all of the docking programs and scoring function combinations.

Thirteen different scoring functions were used to rank multiple conformations generated by DOCK 4.0 and the genetic algorithm GAMBLER by Charifson *et al.* [127]. Using this consensus scoring approach unsurprisingly the hit rates were increased while reducing the number of false positives. The scoring functions which performed consistently well were ChemScore, PLP and DOCK.

Summary and conclusions

An extensive summary of currently available docking methods has been presented. Comparisons suggest that the best algorithm for docking is probably a hybrid of various types of algorithm encompassing novel search and scoring strategies. The most useful docking method will not only perform well, but will be easy to use and parametrise, and sufficiently adaptable such that different functionality may be selected, depending on the number of structures to be docked, the available computational resources, and the complexity of the problem. If the parameters cannot be generated quickly then although the algorithm may be computationally efficient, from a practical point of view it is limited. Conversely, a rapid scoring function may not necessarily be able to model some specific interactions.

Algorithms that use the rigid receptor/flexible ligand approximation are well established and the most successful programs have achieved a success rate of between 70–80%. However, in the few examples where protein flexibility is incorporated into the docking algorithm, it is not clear whether the protein conformational states are sampled extensively. Furthermore, incorporating an 'on-the-fly' solvent model into a docking method is a problem which has only re-

cently been addressed with varying degrees of success. Moreover, although current docking methods show great promise, fast and accurate discrimination between different ligands based on binding affinity, once the binding mode is generated, is still a significant problem.

Acknowledgements

We should like to acknowledge the generous support of AstraZeneca and the University of Southampton. JWE is a Royal Society University Research Fellow.

References

1. Kuntz, I.D. *Science*, 257 (1992) 1078–1082.
2. Lybrand, T.P. *Curr. Opin. Struct. Biol.*, 5 (1995) 224–228.
3. Blaney, J.M. and Dixon, J.S. *Perspect. Drug Discov.*, 1 (1993) 301–319.
4. Abagyan, R. and Totrov, M. *Curr. Opin. Chem. Biol.*, 5 (2001) 375–382.
5. Walters, W.P., Stahl, M.T. and Murcko, M.A. *Drug Discovery Today*, 3 (1998) 160–178.
6. Roe, D.C. and Kuntz, I.D. *J. Comput. Aid. Mol. Des.*, 9 (1995) 269–282.
7. Pearlman, D.A. and Murcko, M.A. *J. Comput. Chem.*, 14 (1993) 1184–1193.
8. Pearlman, D.A. and Murcko, M.A. *J. Med. Chem.* 39 (1996) 1651–1663.
9. Stultz, C.M. and Karplus, M. *Proteins*, 40 (2000) 258–289.
10. Rotstein, S.H. and Murcko, M.A. *J. Comput. Aid. Mol. Des.*, 7 (1993) 23–43.
11. Rotstein, S.H. and Murcko, M.A. *J. Med. Chem.*, 36 (1993) 1700–1710.
12. Moon, J.B. and Howe, W.J. *Proteins*, 11 (1991) 314–328.
13. Eisen, M.B., Wiley, D.C., Karplus, M. and Hubbard, R.E. *Proteins*, 19 (1994) 199–221.
14. Nishibata, Y. and Itai, A. *J. Med. Chem.*, 36 (1993) 2921–2928.
15. Böhm, H.J. *J. Comput. Aid. Mol. Des.*, 6 (1992) 61–78.
16. Gehlhaar, D.K., Moerder, K.E., Zichi, D., Sherman, C.J., Ogden, R.C. and Freer, S.T. *J. Med. Chem.*, 38 (1995) 466–472.
17. Dewitte, R.S., Ishchenko, A.V. and Shakhnovich, E.I. *J. Am. Chem. Soc.*, 119 (1997) 4608–4617.
18. Gillet, V., Johnson, A.P., Mata, P., Sike, S. and Williams, P. *J. Comput. Aid. Mol. Des.*, 7 (1993) 127–153.
19. Welch, W., Ruppert, J. and Jain, A.N. *Chem. Biol.*, 3 (1996) 449–462.
20. Kuntz, I.D., Blaney, J.M., Oatley, S.J., Langridge, R. and Ferrin, T.E. *J. Mol. Biol.*, 161 (1982) 269–288.
21. Shoichet, B.K. and Kuntz, I.D. *Chem. Biol.*, 3 (1996): 151–156.
22. Wüthrich, K., Freyberg, V.B., Weber, C., Wider, G., Traber, R., Widmer, H. and Braun, W. *Science*, 254 (1991) 953–954.
23. Cornell, W.D., Cieplak, P., Bayly, C.I., Gould, I.R., Merz, K.M., Ferguson, D.M., Spellmeyer, D.C., Fox, T., Caldwell, J.W. and Kollman, P.A. *J. Am. Chem. Soc.*, 117 (1995) 5179–5197.

24. Jorgensen, W.L. and Tirado-Rives, J. *J. Am. Chem. Soc.*, 110 (1988) 1657–1666.
25. Brooks, B.R., Bruccoleri, R.E., Olafson, B.D., States, D.J., Swaminathan, S. and Karplus, M. *J. Comput. Chem.*, 4 (1983) 187–217.
26. Eldridge, M.D., Murray, C.W., Auton, T.R., Paolini, G.V. and Mee, R.P. *J. Comput. Aid. Mol. Des.*, 11 (1997) 425–445.
27. Muegge, I. and Martin, Y.C. *J. Med. Chem.*, 42 (1999) 791–804.
28. Jones, G., Willett, P., Glen, R.C., Leach, A.R. and Taylor, R. *J. Mol. Biol.*, 267 (1997) 727–748.
29. Ewing, T.J.A. and Kuntz, I.D. *J. Comput. Chem.*, 18 (1997) 1175–1189.
30. Miller, M., Schneider, J., Sathyanarayana, B.K., Toth, M.V., Marshall, G.R., Clawson, L., Selk, L., Kent, S.B.H. and Wlodawer, A. *Science*, 246 (1989) 1149–1152.
31. Lam, P.Y.S., Jadhav, P.K., Eyermann, C.J., Hodge, C.N., Ru, Y., Bachelier, L.T., Meek, J.L., Otto, M.J., Rayner, M.M., Wong, Y.N., Chang, C.H., Weber, P.C., Jackson, D.A., Sharpe, T.R. and Erickson-Viitanen, S. *Science*, 263 (1994) 380–384.
32. Kollman, P.A. *AMBER 5.0*, University of California, San Francisco, 1996.
33. Mangoni, R., Roccatano, D. and Di Nola, A. *Proteins*, 35 (1999) 153–162.
34. Di Nola, A., Roccatano, D. and Berendsen, H.J.C. *Proteins*, 19 (1994) 174–182.
35. Pak, Y. and Wang, S. *J. Phys. Chem. B*, 104 (2000) 354–359.
36. Nakajima, N., Higo, J., Kidera, A. and Nakamura, H. *Chem. Phys. Lett.*, 278 (1997) 297–301.
37. Carlson, H.A. and McCammon, J.A. *Mol. Pharmacol.*, 57 (2000) 213–218.
38. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H. and Teller, E., *J. Chem. Phys.*, 21 (1953) 1087–1092.
39. Goodsell, D.S. and Olson, A.J. *Proteins*, 8 (1990) 195–202.
40. Morris, G.M., Goodsell, D.S., Huey, R. and Olson, A.J. *J. Comput. Aid. Mol. Des.*, 10 (1996) 293–304.
41. Trosset, J.Y. and Scheraga, H.A. *J. Comput. Chem.*, 20 (1999) 412–427.
42. Trosset, J.Y. and Scheraga, H.A. *J. Comput. Chem.*, 20 (1999) 244–252.
43. Némethy, G., Gibson, K.D., Palmer, K.A., Yoon, C.N., Paterlini, G., Zagari, A., Rumsey, S. and Scheraga, H.A. *J. Phys. Chem.*, 96 (1992) 6472–6484.
44. Apostolakis, J., Plückthun, A. and Caffisch, A. *J. Comput. Chem.*, 19 (1998) 21–37.
45. Caffisch, A., Fischer, S. and Karplus, M. *J. Comput. Chem.*, 18 (1997) 723–743.
46. Davis, M.E., Madura, J.D., Luty, B.A. and McCammon, J.A. *Comput. Phys. Commun.*, 62 (1991) 187–197.
47. Abagyan, R., Totrov, M. and Kuznetsov, D. *J. Comput. Chem.*, 15 (1994) 488–506.
48. Abagyan, R. and Totrov, M. *J. Mol. Biol.*, 235 (1994) 983–1002.
49. Ponder, J.W. and Richards, F.M. *J. Mol. Biol.*, 193 (1987) 775–791.
50. Argos, P. and Abagyan, R. *J. Mol. Biol.*, 225 (1992) 519–532.
51. Totrov, M. and Abagyan, R. *Proteins Suppl.* 1 (1997) 215–220.
52. Liu, M. and Wang, S. *J. Comput. Aid. Mol. Des.*, 13 (1999) 435–451.
53. Rarey, M., Kramer, B., Lengauer, T. and Klebe, G. *J. Mol. Biol.*, 261 (1996) 470–489.
54. Bouzida, D., Rejto, P.A., Arthurs, S., Colson, A.B., Freer, S.T., Gehlhaar, D.K., Larson, V., Luty, B.A., Rose, P.W. and Verkhivker, G.M. *Int. J. Quant. Chem.*, 72 (1999) 73–84.
55. Gehlhaar, D.K., Verkhivker, G.M., Rejto, P.A., Sherman, C.J., Fogel, D.B., Fogel, L.J. and Freer, S.T. *Chem. Biol.*, 2 (1995) 317–324.
56. Hart, T.N. and Read, R.J. *Proteins*, 13 (1992) 206–222.
57. Janin, J. *Prog. Biophys. Molec. Biol.*, 64 (1995) 145–166.
58. McMartin, C. and Bohacek, R.S. *J. Comput. Aid. Mol. Des.*, 11 (1997) 333–344.
59. Accelrys Inc., San Diego, CA 92121.
60. Luty, B.A., Wasserman, Z.R., Stouten, P.F.W., Hodge, C.N., Zacharias, M. and McCammon, J.A. *J. Comput. Chem.*, 16 (1995) 454–464.
61. Stouten, P.F.W., Frommel, C., Nakamura, H. and Sander, C. *Mol. Simulat.*, 10 (1993) 97–120.
62. Schrödinger Inc., San Diego, CA 92122.
63. GOLD 1.2, CCDC, Cambridge, UK, 2001.
64. Morris, G.M., Goodsell, D.S., Halliday, R.S., Huey, R., Hart, W.E., Belew, R.K. and Olson, A.J. *J. Comput. Chem.*, 19 (1998) 1639–1662.
65. Clark, K.P. and Ajay, J. *J. Comput. Chem.*, 16 (1995) 1210–1226.
66. Taylor, J.S. and Burnett, R.M. *Proteins*, 41 (2000) 173–191.
67. Nicholls, A. and Honig, B. *J. Comput. Chem.*, 12 (1991) 435–445.
68. Judson, R.S., Jaeger, E.P. and Treasurywala, A.M. *Theochem.*, 114 (1994) 191–206.
69. Caffisch, A., Niederer, P. and Anliker, M. *Proteins*, 13 (1992) 223–230.
70. Klebe, G. and Mietzner, T. *J. Comput. Aid. Mol. Des.*, 8 (1994) 583–606.
71. Rarey, M., Kramer, B. and Lengauer, T. *J. Comput. Aid. Mol. Des.*, 11 (1997) 369–384.
72. Kabsch, W. *Acta Cryst.*, A32 (1976) 922–923.
73. Böhm, H.J. *J. Comput. Aid. Mol. Des.*, 8 (1994) 243–256.
74. Rarey, M., Kramer, B. and Lengauer, T. *Proteins*, 34 (1999) 17–28.
75. Claussen, H., Buning, C., Rarey, M. and Lengauer, T. *J. Mol. Biol.*, 308 (2001) 377–395.
76. Bron, C. and Kerbosch, J. *Comm. of the A.C.M.*, 16 (1973) 575–576.
77. Brint, A.T. and Willett, P. *J. Chem. Inform. Comput. Sci.*, 27 (1987) 152–158.
78. Grindley, H.M., Artymiuk, P.J., Rice, D.W. and Willett, P. *J. Mol. Biol.*, 229 (1993) 707–721.
79. Meng, E.C., Shoichet, B.K. and Kuntz, I.D. *J. Comput. Chem.*, 13 (1992) 505–524.
80. Makino, S. and Kuntz, I.D. *J. Comput. Chem.*, 18 (1997) 1812–1825.
81. Lorber, D.M. and Shoichet, B.K. *Protein Sci.*, 7 (1998) 938–950.
82. Ewing, T.J.A., Makino, S., Skillman, A.G. and Kuntz, I.D. *J. Comput. Aid. Mol. Des.*, 15 (2001) 411–428.
83. Knegtel, R.M.A., Kuntz, I.D. and Oshiro, C.M. *J. Mol. Biol.*, 266 (1997) 424–440.
84. Still, W.C., Tempczyk, A., Hawley, R.C. and Hendrickson, T. *J. Am. Chem. Soc.*, 112 (1990) 6127–6129.
85. Zou, X.Q., Sun, Y. and Kuntz, I.D. *J. Am. Chem. Soc.*, 121 (1999) 8033–8043.
86. Goodford, P.J. *J. Med. Chem.*, 28 (1985) 849–857.
87. Böhm, H.J. *J. Comput. Aid. Mol. Design*, 6 (1992) 593–606.
88. Mizutani, M.Y., Tomioka, N. and Itai, A. *J. Mol. Biol.*, 243 (1994) 310–326.

89. Nelder, J.A. and Mead, R. *Computer J.*, 7 (1965) 308–313.
90. Weber, P.C., Ohlendorf, D.H., Wendoloski, J.J. and Salemme, F.R. *Science*, 243 (1989) 85–88.
91. Schnecke, V. and Kuhn, L.A. *Perspect. Drug Discov.*, 20 (2000) 171–190.
92. Schnecke, V., Swanson, C.A., Getzoff, E.D., Tainer, J.A. and Kuhn, L.A. *Proteins*, 33 (1998) 74–87.
93. Koehl, P. and Delarue, M. *J. Mol. Biol.*, 239 (1994) 249–275.
94. Schnecke, V. and Kuhn, L.A. *Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology, ISMB99*, AAAI Press, 242–251, 1999.
95. Jiang, F. and Kim, S.H. *J. Mol. Biol.*, 219 (1991) 79–102.
96. Gabb, H.A., Jackson, R.M. and Sternberg, M.J.E. *J. Mol. Biol.*, 272 (1997) 106–120.
97. Sternberg, M.J.E., Gabb, H.A. and Jackson, R.M. *Curr. Opin. Struct. Biol.*, 8 (1998) 250–256.
98. Jackson, R.M., Gabb, H.A. and Sternberg, M.J.E. *J. Mol. Biol.*, 276 (1998) 265–285.
99. Sobolev, V., Wade, R.C., Vriend, G. and Edelman, M. *Proteins*, 25 (1996) 120–129.
100. Sobolev, V., Moallem, T.M., Wade, R.C., Vriend, G. and Edelman, M. *Proteins, Suppl. 1* (1997) 210–214.
101. Burkhard, P., Taylor, P. and Walkinshaw, M.D. *J. Mol. Biol.*, 277 (1998) 449–466.
102. Miller, M.D., Kearsley, S.K., Underwood, D.J. and Sheridan, R.P. *J. Comput. Aid. Mol. Des.*, 8 (1994) 153–174.
103. Shoichet, B.K., Bodian, D.L. and Kuntz, I.D. *J. Comput. Chem.*, 13 (1992) 380–397.
104. Leach, A.R. and Smellie, A.S. *J. Chem. Inform. Comput. Sci.*, 32 (1992) 379–385.
105. Metaphorics LLC, Mission Viejo, CA 92691.
106. Baxter, C.A., Murray, C.W., Clark, D.E., Westhead, D.R. and Eldridge, M.D. *Proteins*, 33 (1998) 367–382.
107. Murray, C.W., Baxter, C.A. and Frenkel, A.D. *J. Comput. Aid. Mol. Des.*, 13 (1999) 547–562.
108. Baxter, C.A., Murray, C.W., Waszkowycz, B., Li, J., Sykes, R.A., Bone, R.G.A., Perkins, T.D.J. and Wylie, W. *J. Chem. Inform. Comput. Sci.*, 40 (2000) 254–262.
109. Perola, E., Xu, K., Kollmeyer, T.M., Kaufmann, S.H., Prendergast, F.G. and Pang, Y.P. *J. Med. Chem.*, 43 (2000) 401–408.
110. Pang, Y.P. and Kozikowski, A.P. *J. Comput. Aid. Mol. Des.*, 8 (1994) 669–681.
111. Taylor, R.D., Jewsbury, P.J. and Essex, J.W. *FDS: Flexible ligand and receptor docking with a continuum solvent model and soft-core energy function*. In Preparation.
112. Jorgensen, W.L., *MCPRO 1.4*, Yale University, New Haven, CT, 1996.
113. Leach, A.R., *J. Mol. Biol.*, 235 (1994) 345–356.
114. Desmet, J., Maeyer, D.M., Hazes, B. and Lasters, I. *Nature*, 356 (1992) 539–542.
115. Schaffer, L. and Verkhivker, G.M. *Proteins*, 33 (1998) 295–310.
116. Weiner, S.J., Kollman, P.A., Case, D.A., Singh, U.C., Ghio, C., Alagona, G., Profeta, S. and Weiner, P. *J. Am. Chem. Soc.*, 106 (1984) 765–784.
117. Wang, J., Kollman, P.A. and Kuntz, I.D. *Proteins*, 36 (1999) 1–19.
118. Price, M.L.P. and Jorgensen, W.L. *J. Am. Chem. Soc.*, 122 (2000) 9455–9466.
119. Westhead, D.R., Clark, D.E. and Murray, C.W. *J. Comput. Aid. Mol. Des.*, 11 (1997) 209–228.
120. Hoffmann, D., Kramer, B., Washio, T., Steinmetzer, T., Rarey, M. and Lengauer, T. *J. Med. Chem.*, 42 (1999) 4422–4433.
121. David, L., Luo, R. and Gilson, M.K. *J. Comput. Aid. Mol. Des.*, 15 (2001) 157–171.
122. Smellie, A., Teig, S.L. and Towbin, P. *J. Comput. Chem.*, 16 (1995) 171–187.
123. Vieth, M., Hirst, J.D., Dominy, B.N., Daigler, H. and Brooks, C.L. *J. Comput. Chem.*, 19 (1998) 1623–1631.
124. Vieth, M., Hirst, J.D., Kolinski, A. and Brooks, C.L. *J. Comput. Chem.*, 19 (1998) 1612–1622.
125. Foreman, K.W., Phillips, A.T., Rosen, J.B. and Dill, K.A. *J. Comput. Chem.*, 20 (1999) 1527–1532.
126. Bissantz, C., Folkers, G. and Rognan, D. *J. Med. Chem.*, 43 (2000) 4759–4767.
127. Charifson, P.S., Corkery, J.J., Murcko, M.A. and Walters, W.P. *J. Med. Chem.*, 42 (1999) 5100–5109.