Symmetry constraints for modelling homo-oligomers

Ludwig Krippahl and Pedro Barahona

NOVA-LINCS, DI, FCT-NOVA,
2829-516 Caparica, Portugal
ludi@fct.unl.pt
pb@fct.unl.pt

Abstract. This paper reports current progress in using constraint propagation techniques to take advantage of the symmetry of homo-oligomeric protein complexes. This is possible because the BiGGER docking algorithm relies on constraint propagation to make protein docking more efficient instead of resorting to the more popular approach based on fast Fourier transforms. This allows us to prune the search for candidate models by imposing constraints derived from the symmetry of each complex, which can be inferred from the number of monomers. This not only extends docking predictions from the interaction of two molecules to n-ary complexes, but also improves the quality of results and computation time. We tested our implementation on a set of 73 dimeric and 22 trimeric complexes, and the overall result was an average reduction in computation time by a factor of approximately 25 and a threefold increase in the number of acceptable models retained during the geometric search space.

Keywords: protein docking; symmetry; constraints

1 Introduction

Protein oligomers are assemblies of protein molecules that come together in a well-defined structure. It has long been recognized that this oligomeric organization plays an important role in protein function [17], with the interaction of these monomers playing a crucial role in determining the kinetics of enzimatic reactions. Several other factors contribute to the evolution of protein oligomers, such as better response to environmental variations, greater availability of interacting sites for substrates and other reaction partners, better feedback response through cooperative binding, among others [6]. Conversely, oligomerization may also play a role in protein evolution, with some oligomeric proteins being thought to have evolved from gene duplication or fusion of their constituent monomers [1]. Oligomerization can also be part of dynamic processes, with some proteins changing their oligomerization state as part of their function or some reaction. For example, proteins involved in programmed cell death (apoptosis) and inflammatory reactions [7,23], precursors to fibrilization of alpha-synuclein during the

onset of Parkinson's disease [5] or ATPases regulating the partition of chromossomal DNA during bacterial replication [15]. The importance and abundance of complexes formed by copies of the same protein justifies an effort to take advantage of the symmetry of such complexes for protein docking simulations.

ClusPro, a well-known protein docking server [11], implements a modification of its standard docking algorithm specifically for homo-oligomers. According to the authors [4], the algorithm begins by generating 20,000 models with the DOT geometric docking algorithm [16], docking two copies of the monomeric protein. The next step is to filter these dimeric models using the ClusPro scoring functions, retaining 2,000 models (500 according to a desolvation score and 1,500 using electrostatics score). Finally, for each retained model the complete oligomer is computed by applying symmetry operations based on the dimer generated by the docking algorithm. Models are then rejected if their symmetry is not compatible with that of an oligomer with the specified number of monomers and symmetry. This approach of *a posteriori* filtering does not take advantage of the constraints imposed by the homo-oligomer symmetry, which could be used to reduce computation time and, more importantly, reduce the number of incorrect models generated during the docking search, from which the correct models must be distinguished by the scoring functions. Thus, the authors report that the docking stage is run in 16 processors and takes less than 30 min whereas the selection algorithm and minimization takes a further 3 hours to predict dimers, longer for higher order oligomers.

The reason for this shortcoming in the most widely used docking programs is the dependency on the fast Fourier transform (FFT) technique to compute the geometrical fit between the interacting proteins. The FFT algorithm is used to speed up the computation of a correlation matrix from two matrices describing the docking partners. This correlation matrix, in turn, identifies the relative positions, for a given orientation, where the surface contact is greater. This technique dates back to the early days of protein-protein docking [9] and has been widely used ever since. However, despite its efficiency, FFT requires that the spatial features of the docking problem be transformed into a frequency domain, via a Fourier transform, where the correlation is computed, which makes FFT methods unsuitable for processing geometric constraints in a way that prunes the search space. In contrast, BiGGER [18] uses constraint propagation in a spatial representation of the interacting molecules in order to improve docking efficiency [12]. Like the FFT methods, BiGGER starts by generating a cubic cell matrix representation of each molecule, using grid cells 1Å wide. However, while FFT methods assign numerical values to each cell and then compute the Fourier transform to obtain the correlation matrix, BiGGER uses two grids for each molecule, for the surface and the core regions, and represents each shape with a two-dimensional array of sets of line segments. This makes for a very compact representation, taking about three orders of magnitude less memory than the FFT matrices, and allows an efficient propagation of the basic docking constraints, which are that the molecular surface must be in contact represented by the overlap of cells in the surface grid but molecules cannot overlap,

which is represented by the overlap of core grid cells. Furthermore, since only the best models are kept (typically five thousand out of thousands of millions of possibilities), the lowest contact score of the current set of models retained serves as a lower bound for the contact score of any acceptable model, which is also used to further prune the search space. This results in computation times similar to FFT with much lower memory requirements. It also preserves the spatial representation in a constraint-based framework that is particularly suitable for implementing additional constraints, such as residue contacts [13] or the symmetry of homo-oligomers, which is the topic of this paper

1.1 Symmetry of homo-oligomers

Homo-oligomers, which are assemblies of several copies of the same protein molecule, are nearly always symmetric [19]. Several factors are suspected to contribute to this, from preventing uncontrolled aggregation of proteins to enzymatic versatility, stability and protection against deleterious mutations. Simulations also suggest that, on average, the binding of the proteins improves with symmetry [2], so it is likely that symmetry is, in itself, favoured by natural selections. At the time of writing, the Protein Data Bank contained 109,457 structures. According to the symmetry search criteria, 61,895 of these are asymmetric, 33,160 have cyclic symmetry, 13,143 have dihedral symmetry and the remaining 1,259 structures fall into the tetrahedral, octahedral, icosahedral or helical symmetry groups. The largest particular symmetry group, by far, is that of the dimers, group C_2, with 27,871 structures.

The cyclical symmetry groups have a single symmetry rotation axis that transforms one monomer into each of the other monomers by a rotation of $360/N$, where N is the number of monomers in the complex. For example, in a C_2 dimer one of the monomers is rotated 180 relative to the other. In a C_3 trimer, each monomer is rotated 120 relative to the adjacent monomers, and so forth. This makes it simple to prune the rotational search space with this constraint and, since the adjacent monomers must be placed relative to each other in such a way that their centres lie on a plane perpendicular to the symmetry axis, this also allows us to prune the translational search space. The use of these constraints to reduce the search space is explained in greater detail in the Methods section.

Although this paper covers only the case of the cyclical symmetry group C_n, the other symmetry groups also have rotational symmetries that can be used to prune the search space in the same way, with slight modifications. For example, the D_n dihedral symmetry group has an n-fold rotational symmetry axis and an additional n 2-fold rotational axes perpendicular to the first one. This requires searching both the $360/n$ and the $360/2$ symmetries to find the best interacting pair of monomers, instead of only the $360/n$ symmetry necessary in the C_n case, but otherwise requires no changes in the base algorithm. The application of this method to dihedral symmetries is currently work in progress, as is the extension to other, less well represented, symmetry groups. Furthermore, although this paper only focuses on homo-oligomers, which are composed of copies of the same structure, the method described here can also be applied to pseudo-symmetric

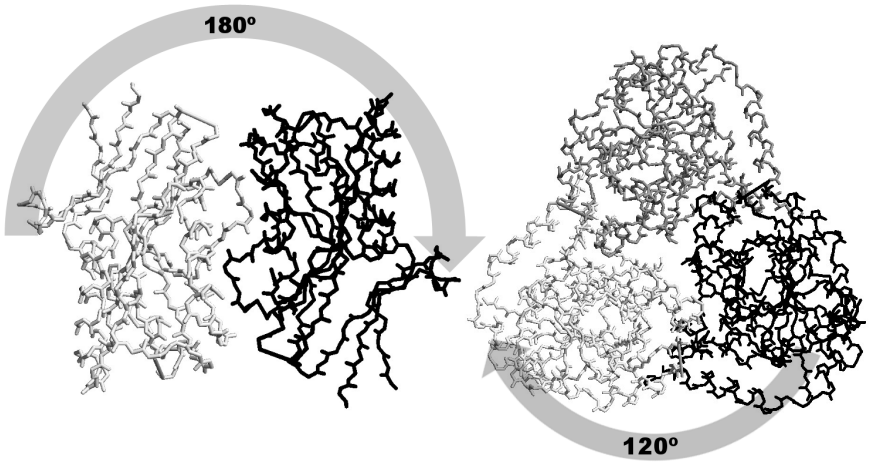

Fig. 1. Two examples of C_n symmetry complexes. On the left, the PyrR attenuation protein, a C_2 dimer (PDB code 1A3C [26]). On the right, the structure of an MTA phosphorylase, a C_3 trimer (PDB code 1CB0 [3])

complexes. These are oligomers in which the monomers are not identical but are similar in structure. The only requirement to apply the method we describe here is that one partner can be structurally aligned with the other partner. This suffices to allow the use of the rotational and translational constraints derived from the symmetry of the complex, requiring only that the partners be structurally aligned before beginning the rotational space search.

2 Method

Docking between two molecules involves a 6 dimensional search over the rotation space of one relative to the other (3 dimensions) and, for each orientation, the translation search of their relative positions (3 dimensions). Symmetry constraints can be used to prune the search space both for the rotations and translations.

For unconstrained docking, BiGGER samples the rotational space by selecting a uniform distribution of axes around which the probe molecule is rotated a number of equally spaced steps. By default, BiGGER generates 289 rotation axes and, for each of these axes, BiGGER generates 23 rotations for an angular search step of 15, giving a total of 6648 orientations (23×289 plus the original orientation). However, since in a C_n complex one monomer must be rotated $360/n$ relative to the previous adjacent monomer, this search can be pruned down to a single rotation of the appropriate angle, $360/n$, around each of a set of rotation axes. Thus we can model the interaction of two consecutive monomers in the cyclical symmetry by sampling only one rotation in each of the same 289 uniformly distributed rotation axes generated for the default,

unconstrained, search. It is important to note that these rotations may or may not be an exact subset of the unconstrained set of rotations depending on the number of monomers on the cyclic symmetry. However, for dimers and trimers, as covered in this paper, the rotations of 180 and 120 respectively are a subset of the unconstrained rotations set.

In addition, for this rotational symmetry to hold, the centre of the rotated monomer must fall into the plane that contains the centre of the reference monomer and is perpendicular to the rotation axis. If this was not the case, then a simple rotation around this axis would not transform one monomer into the other. This constraint allows us to further prune the search space by limiting the translation of the probe to that plane allowed by the symmetry of the complex, as shown in figure 2. Currently, we restrict the translation search to the plane that is perpendicular to the rotation axis and contains the centre of the target protein, plus or minus 3Å.

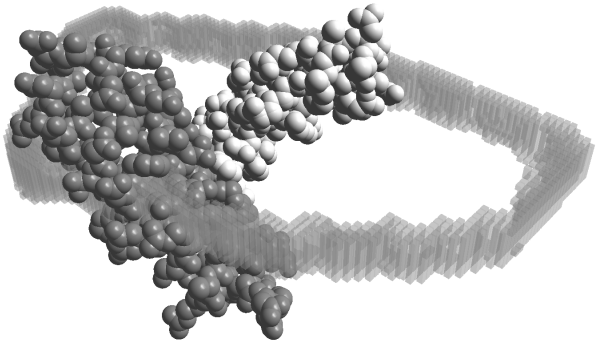

Fig. 2. Representation of the translation domain of the probe molecule (darker shade of grey) relative to the target (lighter shade of grey). The domain is represented as a set of grid cubes, and results from the intersection of the default docking constraints, that the surfaces must be in contact but the cores cannot overlap (see [12]) with the symmetry constraint that the centre of the probe can only be placed on a plane perpendicular to the symmetry axis and containing the centre of the target. The structures represented here are the monomer of a small alpha-helix bundle protein, PDB code 1RPO [27].

To assess the performance gains by applying these constraints we compared the docking of one monomer from each complex of a set of dimeric and trimeric complexes with a copy of itself. Although this is not an ideal test, since the structure is obtained from the complex and not from an unbound protein, it was necessary due to the small number of cases where crystallographers determined the structure of the same protein in both oligomeric and monomeric form. This is the same approach followed by the authors of the ClusPro homo-oligomer docking algorithm [4], precisely for the same reasons. It is important to note, however, that this is not the same as a bound docking simulation, where all partners are

extracted from the bound structure in the right conformation and orientation for a perfect match. In these docking runs, we used only the structure of the first chain, copying that to provide the structure of the partner molecule, exactly as if we only had one monomeric structure available. Though the differences tend to be smaller than they would be if the unbound structure were used, this makes the simulation more realistic than a purely bound docking simulation would be. Furthermore, the contribution of this paper is to show the difference in performance and quality of results that stems from using these symmetry constraints. Some complexes are harder to predict than others, and using unbound structures affects this baseline difficulty, but as the results show the gains are consistent across different complexes.

3 Results

To validate our method, we used the dataset from [21], which focuses on the related problem of inferring the oligomeric state of homo oligomers from crystallographic data. This problem is that the asymmetric unit of the crystal may contain only one monomer even though the structure itself is oligomeric. If this is the case, then the crystal structure determined does not correspond to the structure of the complex, which must be inferred by considering all the possible interfaces allowed by the crystal structure. Though this is not docking, since the interfaces to consider are limited to those present in the crystal lattice and there is no need for the a rotational and translational search, the data required to validate a method for finding the correct oligomeric structure in the crystal is useful for testing a method for finding those structures from the monomers alone. Furthermore, this dataset was also used in the creation of the PISA algorithm for oligomeric inference [14] and the morphological characterization of oligomers in [20]. Symmetric complexes can also be easily obtained from the Protein Data Bank using the symmetry search filters, but this dataset is carefully annotated and selected in order to exclude redundant sequences.

From this dataset, we used 73 of the 77 dimers (1HLR, 1KBA, 5TMP and 1AJQ were rejected due to problems with our automatic chain identification scripts) and the 22 of the 24 trimers (structure 1QEX, for the GP4 protein of bacteriophage T4, is a trimer of dimers and not a trimer of monomers [10] and structure 1CE0 is a small synthetic model for a leucine zipper interaction [22] and thus not a realistic test case). For each complex, we ran three docking runs, in all cases docking the first monomer in the PDB file with a copy of itself. One docking run used no symmetry constraints, one pruned only the rotational search without restricting the translation of the probe relative to the target, and the third docking run for each complex restricted both the rotation and the translation of the probe monomer to those combinations allowed by the symmetry of the complex. The resulting models were classified according to the criteria used in the CAPRI programme [8] to evaluate the quality of protein docking predictions. This classification depends on two values of the root of the mean of the squared distances (rmsd) computed between sets of corresponding atom positions. One

involves superimposing the structure of the target monomer in the model to the corresponding monomer in the known structure of the complex and then measuring the rmsd value between the two positions of the other monomer, the probe, in the model and the known complexes. The other rmsd value is the interface rmsd computed by superimposing the interface atoms of the model to the corresponding interface atoms in the known complex. The interface is defined as the set of residues from one partner that are within 5Å of any residue of the other partner as measured in the known complex. Applying the CAPRI criteria, we considered a model to be acceptable if the ligand rmsd is less than 10Å or the interface rmsd is less than 4Å.

Figure 3 shows the number of acceptable models for the set of dimeric complexes. We chose this measure because the number of acceptable models is the most important factor in this first stage of a docking simulation. The quality of the best models is limited by the resolution of the grid and angular searches and, especially, by the assumption that the proteins remain rigid when interacting. It is only in the final stage of a docking simulation, when a small set of candidates is selected, that one can optimize the structure of the complex with an energy minimization or molecular dynamics simulation, allowing the movement of residue side chains. In contrast, in this first stage we are screening thousands of millions of potential complexes, to retain a few thousand for further scoring and analysis. At this stage the most important goal is to ensure that the largest number of acceptable models is retained, not only to reduce the probability of retaining no acceptable models at this stage and also to aid the subsequent stages, for which the probability of identifying the correct models is increased by having a greater ratio of acceptable models to incorrect candidates.

The complexes were sorted by the number of acceptable models obtained in the unconstrained docking, which serves as the baseline, and the chart compares the number of acceptable models obtained without constraints, with the angular constraint alone and with both the angular and plane constraints. Overall, there is a significant improvement in the number of acceptable models, especially for the cases where few or no acceptable models were obtained in unconstrained docking. In two out of the 73 cases the constrained docking runs result in fewer acceptable models than the unconstrained run, but this seems to be an artefact of the criteria used to classify the models. For example, one of these cases is complex 1DAA, an aminotransferase [24]. In this case the docking with the angular constraint alone did not retain any acceptable models in the set of 5000, but the model closest to being acceptable had an interface rmsd of 4.03Å, which crosses the 4Å threshold by a distance that is on the order of the precision of the rmsd computation. This is not a significant problem and, overall, with the angular constraint alone there is a twofold median increase in the number of acceptable models, rising to threefold when both the angular and plane constraints are enforced.

Although the application of the symmetry constraints results in a significant average improvement, the improvement varies considerably from case to case and, in two cases, actually reducing the number of acceptable models retained.

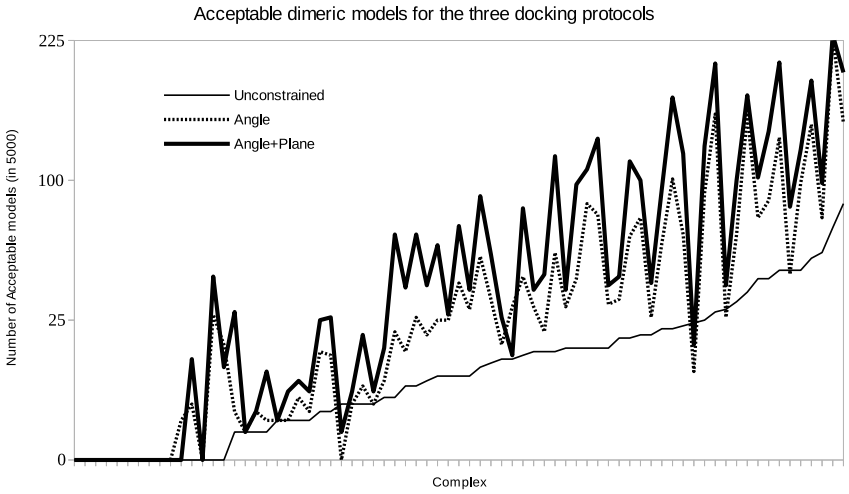

Fig. 3. Comparison of the quality of models obtained for the set of 73 dimeric complexes. The complexes were sorted by the number of acceptable models in the unconstrained dockings (thin black line). Dockings with angular constraints are plotted in a dotted line and dockings with both angular and plane constraints are plotted in a thick black line. Note that the vertical scale is quadratic, not linear, in order to better compare cases for which there were many, few or zero acceptable models.

The reason for this is that, at this first stage of a docking simulation, candidate models must be selected by simple contact scores that can be quickly computed, such as the total contact area, due to the need to screen many millions of possibilities. This is a general problem in protein docking. Thus, the set of models retained are at the extreme of a very long tailed distribution and, even though the search is exhaustive within the limitations of the discrete translation and rotation steps, some acceptable models may be displaced by incorrect models with higher contact scores. Applying symmetry constraints reduces this effect insofar as the incorrect models pruned by the symmetry constraints would have scores higher that acceptable models that would otherwise be left out of the final set. This depends on many details of the protein shapes and the configuration of the target complex, which results in the large variation observed between different cases.

The results for the trimeric complexes, in Figure 4, follow a similar pattern. There was, however, one case worthy of a more detailed discussion. It was the complex with the largest number of acceptable models in the unconstrained docking run, and thus the rightmost complexes in the chart. This is 1AA0, the structure of a viral fibritin [25], a pencil-shaped complex that is very long and oriented along the symmetry axis, as shown in figure 5. This shape makes unconstrained docking easy, resulting in a large number of acceptable models (98)

but makes constrained docking very sensitive to the direction of the rotation axes chosen, since there is only one rotation searched per axis and if the target monomer is not aligned along this axis the target and probe cannot be adequately matched. Fortunately, this is simple to correct by aligning the longest axis of the monomeric structure with the Z axis, for example, before docking. We tested this with a rough manual alignment and obtained 8 acceptable models for this complex. These results are still preliminary and part of current work on optimizing the choice of axes based on the shape of the probe protein, but the important point is that this problem can be easily solved in practice by considering the shape of the monomer to dock when choosing the axes of rotation.

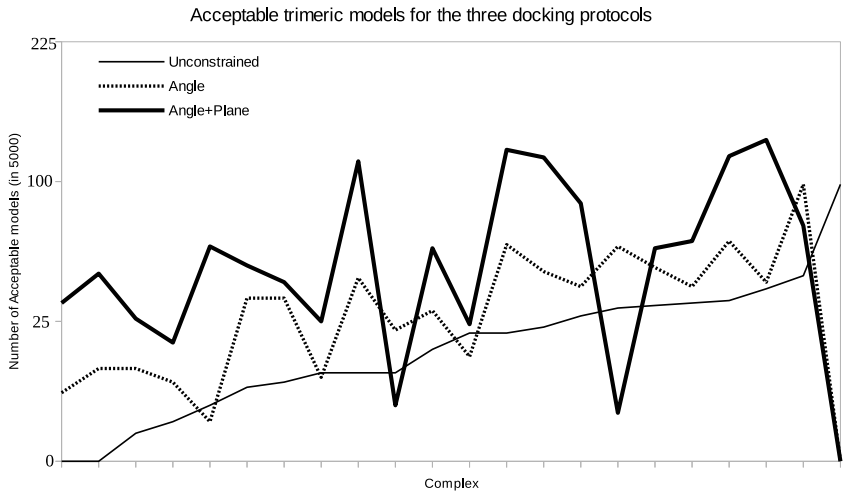

Fig. 4. Comparison of the quality of models obtained for the set of 22 trimeric complexes. The complexes were sorted by the number of acceptable models in the unconstrained dockings (thin black line). Dockings with angular constraints are plotted in a dotted line and dockings with both angular and plane constraints are plotted in a thick black line. Note that the vertical scale is quadratic, not linear, in order to better represent complexes for which there were many, few or zero acceptable models.

Table 1 summarises the results. Using these symmetry constraints reduces computation time by a factor of around 25 and improves the number of acceptable models retained after the geometric search by a factor of three. While the unconstrained running times of BiGGER compare roughly with the geometric search times for ClusPro as reported in [4], computation times for the docking runs that take advantage of the symmetry constraints are much lower, on the order of a few minutes instead of hours. There are some differences between the ClusPro procedure and the one we report here that do not allow for a straightforward comparison of the results. For instance, we are only running the geometric

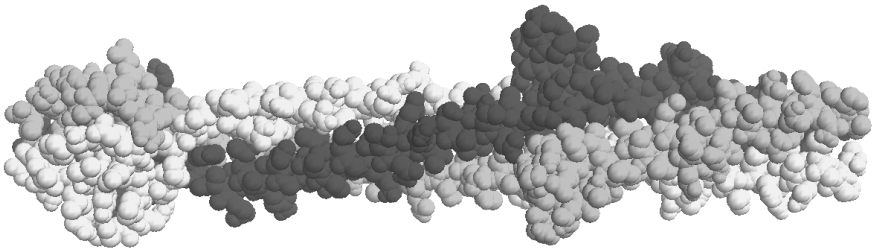

Fig. 5. Representation of the 1AA0 trimer, a long viral fibritin structure [25]. The three monomers are shown in different shades of grey

search stage, without scoring, additional filters or optimization, and we are only retaining 5,000 models instead of the 25,000 reported in [4] for this geometric search stage. However, the main point of this paper is the comparison of docking runs without or without pruning of the search space using symmetry constraints. In this measure, the improvement is significant. In large part this seems to be due to the reduction in the rotational search space (although we did not try using the plane constraint without the angular constraints), but the constraint on the translational search space, difficult to implement on docking algorithms based on FFT, was responsible for an additional 50% improvement on the quality of the results and can result in a significant improvement in computation time once the implementation is optimized for speed, which is not the case in the current prototype. Times reported are for a single core of an AMD FX-8320 CPU running at 3.5 GHz and the memory required for running BiGGER is around 10-15 MB, depending on the structures.

Table 1. This table summarises the results for dimeric and trimeric complexes, showing the computation times, in CPU minutes, and the number of acceptable and good models in each case: unconstrained docking runs (Unconst.); docking runs constrained only with the angular constraints (Angle), and docking runs with both the angular and plane constraints (Plane).

	Dimers			Trimers		
	Unconst.	Angle	Plane	Unconst.	Angle	Plane
Average running time (minutes)	104,3	5,6	4,6	129	5,1	4,2
Median running time (minutes)	74,9	4,1	3,3	103,2	4,7	3,9
Average number of Acceptable models	14,9	38,6	57,4	21,2	32,5	57,6
Median number of Acceptable models	9	25	38	18,5	34	53,5

4 Conclusion and future work

Much of the work reported here is still in progress and there were several issues raised by these results. For example, the angular constraints sometimes reduce the number of acceptable models generated, suggesting that restricting the search to a single angle may not be the best option. We are currently exploring the effect of allowing for some slack in this constraint, including neighbouring angular steps, to account for the sensitivity of docking simulations to small structural and orientation differences. Another parameter to explore is the thickness of the plane in the planar constraint. Currently, we are using a width of 7Å, which allows a slack of ±3Å around the 1Å plane containing the centre of the probe, but this may benefit from optimization. The planar constraint propagation was also not implemented with efficiency in mind, being only a naïve prototype implementation and simple improvements in this code may lead to greater computational efficiency. Finally, the case of the long viral fibritin shows the importance of combining this approach with a good choice of rotation axes, which is also work in progress. However, the overall results clearly show that there is benefit to using the constraints imposed by the symmetry of the complex to prune the search space instead of just filtering the resulting models, as reported in [4]. Not only in the efficiency of the computation, but especially in the quality of the results.

The source code for the implementation of the methods described here is available as part of the Open Chemera Library, at `https://github.com/lkrippahl/Open-Chemera`. The source code is published in the public domain and is free of any copyright restrictions.

5 References

References

1. Ali, M.H., Imperiali, B.: Protein oligomerization: how and why. Bioorganic & medicinal chemistry 13(17), 5013–5020 (2005)
2. André, I., Strauss, C.E., Kaplan, D.B., Bradley, P., Baker, D.: Emergence of symmetry in homooligomeric biological assemblies. Proceedings of the National Academy of Sciences 105(42), 16148–16152 (2008)
3. Appleby, T.C., Erion, M.D., Ealick, S.E.: The structure of human 5-deoxy-5-methylthioadenosine phosphorylase at 1.7 å resolution provides insights into substrate binding and catalysis. Structure 7(6), 629–641 (1999)
4. Comeau, S.R., Camacho, C.J.: Predicting oligomeric assemblies: N-mers a primer. Journal of structural biology 150(3), 233–244 (2005)
5. Conway, K.A., Lee, S.J., Rochet, J.C., Ding, T.T., Williamson, R.E., Lansbury, P.T.: Acceleration of oligomerization, not fibrillization, is a shared property of both α-synuclein mutations linked to early-onset parkinson's disease: implications for pathogenesis and therapy. Proceedings of the National Academy of Sciences 97(2), 571–576 (2000)
6. D'Alessio, G.: The evolutionary transition from monomeric to oligomeric proteins: tools, the environment, hypotheses. Progress in biophysics and molecular biology 72(3), 271–298 (1999)

7. Eskes, R., Desagher, S., Antonsson, B., Martinou, J.C.: Bid induces the oligomerization and insertion of bax into the outer mitochondrial membrane. Molecular and cellular biology 20(3), 929–935 (2000)
8. Janin, J.: Assessing predictions of protein-protein interaction: the capri experiment. Protein Sci 14(2), 278–283 (Feb 2005)
9. Katchalski-Katzir, E., Shariv, I., Eisenstein, M., Friesem, A.A., Aflalo, C., Vakser, I.A.: Molecular surface recognition: determination of geometric fit between proteins and their ligands by correlation techniques. Proceedings of the National Academy of Sciences 89(6), 2195–2199 (1992)
10. Kostyuchenko, V.A., Navruzbekov, G.A., Kurochkina, L.P., Strelkov, S.V., Mesyanzhinov, V.V., Rossmann, M.G.: The structure of bacteriophage t4 gene product 9: the trigger for tail contraction. Structure 7(10), 1213–1222 (1999)
11. Kozakov, D., Beglov, D., Bohnuud, T., Mottarella, S.E., Xia, B., Hall, D.R., Vajda, S.: How good is automated protein docking? Proteins: Structure, Function, and Bioinformatics 81(12), 2159–2166 (2013)
12. Krippahl, L., Barahona, P.: Applying constraint programming to rigid body protein docking. In: Beek, P. (ed.) Principles and Practice of Constraint Programming - CP 2005. Lecture Notes in Computer Science, vol. 3709, pp. 373–387. Springer, Berlin Heidelberg (2005)
13. Krippahl, L., Barahona, P.: Protein docking with predicted constraints. Algorithms for Molecular Biology 10(1), 9 (2015)
14. Krissinel, E., Henrick, K.: Inference of macromolecular assemblies from crystalline state. Journal of molecular biology 372(3), 774–797 (2007)
15. Leonard, T.A., Butler, P.J., Löwe, J.: Bacterial chromosome segregation: structure and dna binding of the soj dimera conserved biological switch. The EMBO journal 24(2), 270–282 (2005)
16. Mandell, J.G., Roberts, V.A., Pique, M.E., Kotlovyi, V., Mitchell, J.C., Nelson, E., Tsigelny, I., Ten Eyck, L.F.: Protein docking using continuum electrostatics and geometric fit. Protein Eng 14(2), 105–113 (Feb 2001)
17. Monod, J., Wyman, J., Changeux, J.P.: On the nature of allosteric transitions: a plausible model. Journal of molecular biology 12(1), 88–118 (1965)
18. Palma, P.N., Krippahl, L., Wampler, J.E., Moura, J.J.: Bigger: a new (soft) docking algorithm for predicting protein interactions. Proteins 39(4), 372–384 (Jun 2000)
19. Plaxco, K.W., Gross, M.: Protein complexes: the evolution of symmetry. Current Biology 19(1), R25–R26 (2009)
20. Ponstingl, H., Kabir, T., Gorse, D., Thornton, J.M.: Morphological aspects of oligomeric protein structures. Progress in biophysics and molecular biology 89(1), 9–35 (2005)
21. Ponstingl, H., Kabir, T., Thornton, J.M.: Automatic inference of protein quaternary structure from crystals. Journal of Applied Crystallography 36(5), 1116–1122 (2003)
22. Shu, W., Ji, H., Lu, M.: Trimerization specificity in hiv-1 gp41: analysis with a gcn4 leucine zipper model. Biochemistry 38(17), 5378–5385 (1999)
23. Srinivasula, S.M., Ahmad, M., Fernandes-Alnemri, T., Alnemri, E.S.: Autoactivation of procaspase-9 by apaf-1-mediated oligomerization. Molecular cell 1(7), 949–957 (1998)
24. Sugio, S., Petsko, G.A., Manning, J.M., Soda, K., Ringe, D.: Crystal structure of a d-amino acid aminotransferase: how the protein controls stereoselectivity. Biochemistry 34(30), 9661–9669 (1995)

25. Tao, Y., Strelkov, S.V., Mesyanzhinov, V.V., Rossmann, M.G.: Structure of bacteriophage t4 fibritin: a segmented coiled coil and the role of the c-terminal domain. Structure 5(6), 789–798 (1997)
26. Tomchick, D.R., Turner, R.J., Switzer, R.L., Smith, J.L.: Adaptation of an enzyme to regulatory function: structure of bacillus subtilis pyrr, a pyr rna-binding attenuation protein and uracil phosphoribosyltransferase. Structure 6(3), 337–350 (1998)
27. Vlassi, M., Steif, C., Weber, P., Tsernoglou, D., Wilson, K.S., Hinz, H.J., Kokkinidis, M.: Restored heptad pattern continuity does not alter the folding of a four-α-helix bundle. Nature Structural & Molecular Biology 1(10), 706–716 (1994)