

Workshop #7: Docking

Protein-protein docking is the prediction of a complex structure starting from its monomer components. The search space can be extremely large, so large amounts of computational resources are typically required. In this workshop, we will explore several of the techniques briefly; keep in mind that for real applications, many more decoys will need to be tested.

Suggested Readings

1. J. J. Gray *et al.*, “Protein-protein docking with simultaneous optimization of rigid-body displacement and side-chain conformations,” *J. Mol. Biol.* **331**, 281-299 (2003).
2. S. Chaudhury & J. J. Gray, “Conformer selection and induced fit in flexible backbone protein-protein docking using computational and NMR ensembles,” *J. Mol. Biol.* **381**, 1068-1087 (2008).

Fast-Fourier transform based docking via ZDOCK

There are several servers available based on fast Fourier transforms. These servers are able to quickly carryout a global, grid-based matching search.

1. Go to the ZDOCK server (<http://zdock.bu.edu>) and upload trypsin (2PTN) and its inhibitor (1BA7 chain B) for docking. If completing this workshop for a class, do this in groups in order to not overload the server. When the jobs have finished (typically under an hour), download the top five models. Are these models similar or diverse? How so?
2. Are any of the models similar to the crystal structure of the bound complex (1AVW)?

Other servers include SmoothDock (<http://structure.pitt.edu/servers/smoothdock>), ClusPro (<http://nrc.bu.edu/cluster>), Haddock (<http://haddock.chem.uu.nl>), and GRAMM-X (<http://vakser.bioinformatics.ku.edu/resources/gramm/grammx>). Any of these provide global docking services to create models which might be useful for refinement by RosettaDock.

Docking moves in Rosetta

For the following exercises, you may use either the bound complex of trypsin or the unbound components. To use the unbound components, you will first need to make a pdb file which has coordinates of both chains (e.g., use the Linux command `cat` or use a text editor).

The fundamental docking move is a rigid-body transformation consisting of a translation and rotation. Any rigid body move also needs to know which part moves and which part is fixed. In

Rosetta, this division is known as a “jump” and the set of protein segments and jumps are stored in an object called a “fold tree.” These objects are set up for you by the DockingProtocol object:

```
DockingProtocol().setup_foldtree(pose)
DockingProtocol().setup_foldtree(pose, "A_B")
print pose.fold_tree()
```

3. In the fold tree printout, each three number sequence is the beginning and ending residue number, then a code. The codes are “-1” for stretches of protein and “1” for a jump. How many jumps are there in your pose?

You can see the type of information in the jump by printing it from the pose:

```
jump_num = 1
print pose.jump(jump_num)
```

4. Write out the rotation matrix and the translation vector defined by the jump.

The two basic manipulations are translations and rotations. For translation, the change in x , y , and z coordinates are needed as well as the jump number. A rotation requires a center and an axis about which to rotate. The rigid-body displacement can be altered directly with the RigidBodyTransMover for translations or the RigidBodySpinMover for rotations.

For structure prediction calculations, we have a Mover which is preconfigured to make random movements of set magnitudes (in this case, 8° rotation and 3 Å translation):

```
pert_mover = RigidBodyPerturbMover(jump_num, 3, 8)
```

5. Apply the RigidBodyPerturbMover to a pose and output the structure to a file. Load the structure into PyMol to confirm the motions are what you expect. What are the new rotation matrix and translation vector in the jump? How many Ångstroms did the downstream protein move?

Global perturbations are useful for making completely randomized starting structures. The following mover will rotate a protein about its geometric center. The final orientation is equally distributed over the ‘globe’.

```
randomize1 = RigidBodyRandomizeMover(pose, jump_num,
    partner_upstream)
randomize2 = RigidBodyRandomizeMover(pose, jump_num,
    partner_downstream)
```

6. Apply both movers to the starting structure, and view the structure in PyMol (you might view it along with the original pose). Does the new conformation look like a candidate docked structure yet?

Since proteins are not spherical, sometimes the random orientation creates severe clashes between the docking partners, and other times it places the partners so they are not longer touching. The DockingSlideIntoContactMover will translate the downstream protein along the line of protein centers until the proteins are in contact.

```
slide = DockingSlideIntoContact(jump_num)
slide.apply(pose)
```

The MinMover, which we have previously used to change torsion angles to find the nearest minimum in the score function, can also operate on the jump translation and rotation. It suffices to set the jump variable as moveable in the MoveMap:

```
movemap = MoveMap()
movemap.set_jump(jump_num, True)

minmover = MinMover()
minmover.movemap(movemap)
minmover.score_function(scorefxn) # use any scorefxn
minmover.apply(pose)
```

7. Apply the MinMover. How much does the score change? What are the new rotation matrix and translation vector in the jump? How many Ångstroms did the downstream protein move?

Low-resolution docking via RosettaDock

RosettaDock can also perform global docking runs, but it can require significant time. Typically, 10^5 to 10^6 decoys are needed in a global run. For this workshop, we will create a much smaller number and learn the tools needed to handle large runs.

Docking is available as a Mover which completely encompasses the protocol. To use the mover, you will need a starting pose with both chains and a jump defined. The structure must be in low-resolution (centroid) mode (use the SwitchResidueTypeSetMover), and you will need a low-resolution score function:

```
scorefxn_low = create_score_function('interchain_cen')
```

Create low-resolution structures as follows:

```
dock_lowres = DockingLowRes(scorefxn_low, jump_num)
dock_lowres.apply(pose)
```

8. You can compare structures by calculating the root-mean-squared deviation of all the C_{α} atoms, using `CA_rmsd(pose1, pose2)`. In docking, a more useful measure is the ligand RMSD, which is the deviation of the backbone C_{α} atoms of the ligand after superposition of the receptor protein backbones. You can calculate ligand RMSD with `DockingProtocol().calc_Lrmsd(pose1, pose2)`. Using both measures, how far did your pose move from the low-resolution search?

9. Examine the decoy created in PyMol. Does it look like a reasonable structure for a protein-protein complex? Explain.

Job distributor

For exhaustive searches with Rosetta (docking, refinement, or folding), it is necessary to create a large number of candidate structures, termed decoys. This is often accomplished by spreading out the work over a large number of computers. Additionally, each decoy created needs to be individually labeled. The object which is responsible for managing the output is called a Job Distributor. Here, we will use a simple job distributor to create multiple structures. The following constructor sets the job distributor to create 10 decoys, with filenames like `output1.pdb`, `output2.pdb`, etc. The `pdb` files will also include scores according to the `ScoreFunction` provided.

```
jd = JobDistributor("output", 10, scorefxn_low)
```

It is also useful to compare each decoy to the native structure (if it is known, otherwise any reference structure can be used). The job distributor will do the `rmsd` calculation and final scoring upon output. Set the native pose:

```
native_pose = Pose("your_favorite_protein.pdb")
jd.native_pose = native_pose
```

Now, run the low-resolution docking protocol to create a structure, and output a decoy:

```
dock_lowres.apply(pose)
jd.output_decoy(pose)
```

Do this twice and confirm that you have two output files.

We can finish the set by using the `JobDistributor` to set up a loop:

```
while (jd.job_complete == False):
    pose.assign(starting_pose)
    dock_lowres.apply(pose)
    jd.output_decoy(pose)
```

Note the `jd.job_complete` Boolean that indicates whether all 10 decoys have been created.

10. Run the loop to create 10 structures. The job distributor also creates a score file, `output.sc`, which summarizes the energies and rmsds of all structures created. Examine that file. What is the lowest score? What is the lowest energy?
11. Reset the JobDistributor to create 100 decoys (or more or less, as the speed of your processor allows) by reconstructing it. Rerun the loop above to make 100 decoys. Use your score file to plot score versus rmsd. Two easy ways to do this are to import the score file into Excel, or to use the linux command `gnuplot`. Do you see a funnel?

High-resolution docking

The high-resolution stage of RosettaDock is also available as a Mover. This mover encompasses random rigid-body moves, side-chain packing, and gradient-based minimization in the rigid-body coordinates. High-resolution docking needs an all-atom score function. The optimized docking weights are available as a patch to the standard all-atom energy function.

```
scorefxn_high = create_score_function_ws_patch('standard', 'docking')
dock_hires = DockingHighRes(scorefxn_high, jump_num)
```

A high-resolution decoy needs side chains. One way to place the side chains is to call the `PackMover`, which will generate a conformation from rotamers. A second way is to copy the side chains from the original monomer structures. This is often helpful for docking calculations since the monomer crystal structures have good side chain positions.

```
recover_sidechains = ReturnSidechainMover(starting_pose)
recover_sidechains.apply(pose)
```

12. Load one of your low-resolution decoys, add the side chains from the starting pose, and refine the decoy using high-resolution docking. How far did the structure move during refinement? How much did the score improve?

Starting from your lowest-scoring low-resolution decoy, create three high-resolution decoys (you might use the JobDistributor). Do the same starting from the native structure.

13. How do the refined-native scores compare to the refined-decoy scores?

14. What is the rmsd of the refined native? Why is it not zero?

15. How much variation do you see in the refined native scores? In the refined decoy scores? Is the difference between the refined natives and the refined decoys significant?

Docking Funnel

16. Using a job distributor, create 10 decoys starting with a `randomize2` perturbation, 10 decoys starting from different local random perturbations (8° , 3 Å), 10 decoys starting from low-resolution decoys, and 10 starting from the native structure. Plot all of these points on a funnel plot. How is the sampling from each method? Does the scoring function discriminate good complexes?

Conformer Selection for Ensemble Docking

Ensemble docking can use multiple backbones for one or both docking partners. One application is the use of NMR structures for docking. NMR pdb files include multiple models (typically 30-50), all of which are reasonable solutions to the spectroscopy constraints.

17. Load the NMR file of 1EGL. How many models are in this structure?

During docking, conformers can be changed using the `ConformerSwitchMover`. Construct the mover with the ensemble of backbones:

```
conformerswitch = ConformerSwitchMover(239,301,jump_num,  
scorefxn_low,"1EGL.pdb")
```

18. Apply the conformer selection mover and confirm that the backbone changed by inspecting the ϕ and ψ angles. Write down a pair of old and new (ϕ,ψ) values and the residue number.

19. Write a conformer selection docking procedure which alternates between `RigidBodyPerturbMoves` and `ConformerSwitchMoves`. Loop through both moves 50 times. Run this code starting from the native structure. Is the final backbone selected close to the native bound structure?

Induced-Fit Docking

In induced-fit binding, a protein changes conformation due to interactions with the partner. We can emulate this simply by allowing gradient-based minimization along the backbone torsion angles.

You can use the following syntax to use your own MoveMap in the DockingHighRes mover:

```
dock_hires_flex = DockingHighRes(scorefxn_high, jump_num)
movemap = MoveMap()
dock_hires_flex.movemap(movemap)
dock_hires_flex.apply(pose)
```

How would you configure a MinMover to vary both backbone torsions and the docking rigid-body displacement?

Programming assignments

1. Output a structure with a 10 Å translation, and another with a 30° rotation (both starting from the same starting structure), and load them into PyMol to confirm the motions are what you expect.
2. *Diffusion.* Make a series of random rigid body perturbations, and record the rmsd after each. Plot rmsd versus the number of moves. Does this process emulate diffusion? If it did, how would you know? (Hint: there is a way to plot these data to make them linear.)
3. Create 10 structures using the ClassicRelax protocol and use those structures for docking. Do you get better results using an ensemble of relaxed crystal structures, or using an ensemble from NMR?
4. Starting from a low-resolution decoy, refine the structure in three separate ways:
 - a. side-chain packing
 - b. gradient-based minimization in the rigid-body coordinates
 - c. gradient-based minimization in the torsional coordinates
 - d. the docking high-resolution protocolFor each, note the change in rmsd and the change in score. Which operations move the protein the most? Which make the most difference in the score?
5. Using the MonteCarlo object, the RigidBodyMover, PackRotamers, and the MinMover, create your own high-resolution docking protocol.
Bonus: Can you tune it to beat the standard protocol? “Beating” the standard protocol could mean achieving lower energies, running in faster time, and/or being able to better predict complexes.