

Workshop #6: Packing & Design

Suggested readings

1. J. Desmet *et al.*, “The dead-end elimination theorem and its use in protein side-chain positioning,” *Nature* **356**, 539-543 (1992).
2. B. Kuhlman & D. Baker, “Native protein structures are close to optimal for their structures,” *PNAS* **97**, 10383, 2000.

Rosetta uses a Monte Carlo optimization routine to pack side chains using a library of conformations, or rotamers. This operation can be used for side-chain packing for operations like refinement or for designing optimal sequences. This workshop will examine both capabilities.

Side-chain conformations, the rotamer library, and Dunbrack energies

Begin by loading cetuximab from PDB 1YY8.

1. What are the ϕ , ψ , and χ angles of residue K49?
2. Open the file `lys.*.bbdep.regular.lib` from the `minirosetta_database/dun??` directory. Find the ϕ/ψ bin for the K49 position and find the nearest rotamer. What are the χ angles and standard deviations of this rotamer? What is its probability?
3. Score your pose with the standard full-atom score function. What is the energy of K49? Note the Dunbrack energy component (`fa_dun`), which represents the side-chain conformational probability. Does it match that which you found in the table (you will need to convert probability to energy; use $kT = 1$)? If not, why not?
4. Use `pose.set_chi(i, j, chi)` to set the side chain of residue 49 to the all-trans conformation (here, i is the χ index, j is the residue number, and chi is the new torsion angle in degrees). Re-score the pose and note the Dunbrack energy. Does it match the probability in the table? Is this conformation valid for Cetuximab (*i.e.*, is the total score of this residue acceptable)?

Monte Carlo side-chain packing

Side-chain packing can be done in a Monte Carlo search routine which iteratively swaps rotamers of a random residue and tests each move using the Metropolis criterion. Rosetta has such a routine pre-packaged as a mover which carries out a simulated annealing search each time it is applied. The specific scope of the packing is specified in a PackerTask object, which we can specify via commands or from an input file. Create a PackerTask as follows. This will set the task to allow packing only of residue 49:

```
task_pack = standard_packer_task(pose)
task_pack.restrict_to_repacking()
task_pack.fix_everything()
task_pack.set_pack_residue(49, True)
```

Confirm your settings using

```
print task_pack
```

We now can create a PackRotamersMover:

```
packmover = PackRotamersMover(scorefxn, task_pack)
```

5. Apply the packmover to your pose with `packmover.apply(pose)`. Now what are the χ angles of K49? Which rotamer is this? What is the Dunbrack energy?
6. What is the new total energy of K49? Why did Rosetta pick this rotamer? Answer this in terms of the components of the score function and in terms of the residues with which K49 interacts.

Packing for refinement

Side-chain packing can be used when converting a pose from centroid to full-atom mode, and it is used extensively in full-atom refinement calculations. Let's examine how packing improves scores.

Use your code from last week to create a centroid-representation model for a zinc-finger sequence (or use one of your structures from the last workshop). Save that centroid "decoy" so that we can compare several basic refinement steps.

7. Load the centroid decoy and convert it to full-atom representation using the `SwitchResidueTypeSetMover`. Save this starting configuration for future use. Score the pose. Why is the score so high?

8. Create a default `PackRotamersMover` with a `PackerTask` that allows all residues to vary χ angles. Create a test pose from your start pose and pack the side chains. What is the new pose score?

9. Reset the test pose to the start configuration. Create a `MinMover` using the ‘dfpmin’ minimization scheme. Create a `MoveMap` that allows χ angles but *not* $\phi/\psi/\omega$ angles to vary. Apply the `MinMover` and rescore the pose. How does this energy compare?

10. Again reset the test pose. This time convert the pose to full-atom representation, apply the packer, and then minimize on the χ angles. Now what is the final score?

For fun, you might examine the individual residue energies to find the residues most responsible for the score changes. Typically, a small number of residues may make clashes which can be resolved using the χ angle minimization, which allows off-rotamer side-chain conformations.

Design

Design calculations can be accomplished simply by packing side chains with a rotamer set that includes all amino acid types. That is, when the Monte Carlo routine swaps rotamers, it could replace the exiting side chain with another conformation of the same residue, or some conformation of a different residue type. Trial mutations are accepted with the Metropolis criterion, and the standard full-atom energy function is supplemented by a reference energy term which represents the relative energies of each residue type in an unfolded peptide.

Design operations are easiest to specify through a data file called a “resfile.” You can create a resfile for a given pdb file or pose using:

```
generate_resfile_from_pdb("1YY8.pdb", "1YY8.resfile")
generate_resfile_from_pose(pose, "1YY8.resfile")
```

Inside the resfile you will see a list of all residues and NATRO next to it, indicating that it is set to use the native rotamer. NATRO can be changed to the following:

NATRO	use native amino acid and native rotamer (does not repack)
NATAA	use native amino acid, but allow repacking to other rotamers
PIKAA ILV	use only the following amino acids and allow repacking between them
ALLAA	use all amino acids and all repacking

Edit the resfile to allow force residue 49 to be glutamic acid (“49 A PIKAA E”) and save the file as 1YY8-K49E.resfile. Create a new task for design from the resfile:

```
task_design = TaskFactory.standard_packer_task(pose)
task_design.read_resfile("1YY8-K49E.resfile")
```

- Score the original conformation from the pdb for reference. Create a new PackResiduesMover with the design task and use it to mutate residue 49 to glutamic acid. What is the predicted ΔG of the mutation? Is this a stabilizing mutation?
- Note the residue reference energy term (“ref”) in the scoring function. What is this value before and after you mutated the residue? What does this energy represent?
- Create a new PackerTask and PackRotamersMover using a new resfile which allows residue 49 to be designed freely (“49 A ALLAA”), and apply the mover. What residue does Rosetta choose? Why?
- Create your own resfile which will restrict residue 49 to only negatively charged residues using the resfile line “49 A PIKAA DE” and re-apply the design mover. Now what residue is chosen? What is the new residue energy, and why (physically) is it less favorable than the last design?
- Let’s try to make this design more favorable. Select several surrounding residues for design, and set them also to enable mutations to any residue. Call the design mover again. Now what do you find?

Programming exercises

- Refinement and discrimination.* Download the “single misfold” decoy set from the Decoys ‘R Us repository at dd.compbio.washington.edu/ddownload.cgi?misfold.

(Documentation for this project is at dd.compbio.washington.edu.) This repository has a single “correct” and “incorrect” predicted structure for several proteins. For this exercise, analyze pDBs 2ci2 and 2cro; each has two “incorrect” structures offered. (Technical note: These decoys have an empty Occupancy field in the PDB ATOM records; a value of 1 needs to be added before Rosetta will load these structures.)

Write a program which will calculate and output the score for each decoy (i) as is from the PDB file (ii) after packing only (iii) after minimization only (iv) after packing and minimizing. For each of the four cases, compare the scores of the “correct” structure with those of the “incorrect” structure. Which schemes successfully discriminate the correct structures?

2. Write a refinement protocol which will iterate between side-chain packing, small and shear moves, and minimization. Where is the best place to position the Monte Carlo acceptance test? Test your protocol by making 10 independently-refined structures for the correct and incorrect decoys of 2cro from the Decoys ‘R Us single misfold set. Is this protocol able to discriminate the correct decoy? Submit your code.
3. HIV-1 protease is a major drug target for antiretroviral therapies. Protease inhibitors are designed from substrate peptide mimics. We will attempt to take a natural substrate peptide of HIV-1 protease and design it for improved binding—potentially to serve as a good template for drug design. Use PDB file 1kjg for the following analysis.
 - a. Turn on side-chain packing for the protease active site (residues 8, 23, 25, 29, 30, 32, 45, 47, 50, 53, 82 and 84 of both chain A and B) and for the peptide (residues 2-9 on chain P; all of these numbers follow the PDB numbering).
 - b. Repack the above side chains and then energy minimize those same side chains with the backbone fixed. Generate 10 decoys and record the energies for each decoy. This will represent the reference state: the wild-type peptide bound to the protease.
 - c. For residue 2 of the peptide (chain p) allow repacking to any of the 20 amino acids, while leaving the packing and side-chain minimization the same as in step 3. Generate 10 decoys and record the energies. These will represent single mutants at that residue position.
 - d. Repeat step c for each of the other 8 residues in the substrate peptide.
 - e. Take the lowest energy for each mutation position and compare it the lowest energy for the wild type. Do single mutants at any of these positions improve the energy over the wild type? Which ones? By how much? Which energy components are mostly responsible?
 - f. What peptide residue positions are easiest to improve on? Which positions are the hardest?
 - g. Are there any other trends? Hydrophobic vs. polar, bulky residues vs. small residues, etc.
 - h. Altman *et al.* (*Proteins* 2008) found, using their own computational design algorithm, that the most favorable sequences were a triple mutant E3D/T4I/V6L, a single mutant T4V, and a single mutant D3Q. How do their results compare with yours?

- i. Natural substrates are often sub-optimal binders. Why would this be advantageous?
4. *Effect of backbone conformation on design.* HIV-1 protease is promiscuous, meaning it can cleave a wide range of peptides beyond the ten natural substrates of the virus. Let's examine the preferences of the enzyme through Rosetta design calculations.
 - a. Download the complex of HIV-1 protease in complex with CA-P2 peptide (1f7a). Select the eight peptide residues for unrestricted design and let Rosetta redesign the substrate sequence. What is the new sequence and how does it compare to the original? What percent of the original sequence was optimal for its structure?
 - b. Download the complex of HIV-1 protease in complex with RT-RH peptide (1kjg). Note that the enzyme is the same here, but it is crystallized with a different substrate. Again, design the eight substrate residues with Rosetta. What percent of this substrate sequence is optimal for this crystal structure?
 - c. How do the designed sequences of (a) and (b) compare? Why should they be the same? Why would they not be the same? What are the implications for the field of computational protein design?
5. Write a program which iterates between design of all residues of a protein and refinement via small, shear, and minimization moves.

Thought questions

1. Amino acid reference energy
 - a. What is the thermodynamic meaning of the `ref` energy term, and what does it correspond to *physically*?
 - b. During evolution, the genome sequence may mutate to cause protein sequence changes. Alternately, one could consider the difference in evolutionary propensities for each residue type. How could you derive reference energies from sequence data, and what would that mean?
 - c. How do Kuhlman & Baker fit the reference energies in their 2000 PNAS paper?

References

1. S. C. Lovell *et al.*, "The penultimate rotamer library," *Proteins* **40**, 389-408 (2000).
2. R. L. Dunbrack & F. E. Cohen, "Bayesian statistical analysis of protein side-chain rotamer preferences," *Protein Sci.* **6**, 1661-1681 (1997).