

Ligand Docking

Bold text means that these files and/or this information is provided.

Italicized text means that this material will NOT be conducted during the workshop

fixed width text means you should type the command into your terminal

If you want to try making files that already exist (e.g., input files), write them to a different directory! (mkdir my_dir)

In addition to following this sample docking problem, the user is encouraged to review the Rosetta user guide including the section on ligand-centric movers for use with RosettaScripts.

<https://www.rosettacommons.org/docs/latest/>

Overview

This small-molecule docking tutorial consists of 3 parts: standard ligand docking, high-throughput screening and RosettaLigandEnsemble.

1. Standard ligand docking: **1_vanilla_docking/**

Determining the binding conformation of a small-molecule ligand within a pre-defined pocket.

2. High-Throughput Screening: **2_vHTS/**

This tutorial will walk through the setup process for vHTS using RosettaLigand. This is generally better for medium-sized libraries (<50,000 compounds). In this tutorial, we will only be running this for a handful of ligands, but the same process can be applied to a much larger set.

3. RosettaLigand Ensemble: **3_Ensemble_docking/**

RLE was developed based on the assumption that similar ligands bind in a similar manner. The set of ligands is superposed and subsequently docked. This is meant to be run in conjunction with experimental SAR data commonly acquired in medicinal chemistry campaigns.

Due to time, I would highly suggest to go through the standard ligand docking procedure first as an introduction to how Rosetta handles small molecules then moving to either HTS or Ensemble docking depending on your specific interests.

Ligand Docking with a G-Protein Coupled Receptor

The experimental data for this tutorial is derived from: **Chien, E. Y. T. et al. Structure of the human dopamine D3 receptor in complex with a D2/D3 selective antagonist. Science 330, 1091-5 (2010).**

This particular D3/eticlopride protein-ligand complex was used as a target in the GPCR Dock 2010 assessment, the results of which are discussed here: **Kufareva, I. et al. Status of GPCR modeling and docking as reflected by community-wide GPCR Dock 2010 assessment. Structure 19, 1108-1126 (2011).**

If you are interested in more information on the performance of Rosetta in modeling and docking D3/GPCRs in general, please consult **Nguyen, E. D. et al. Assessment and challenges of ligand docking into comparative models of g-protein coupled receptors. PLoS One 8, (2013).**

Dopamine is an essential neurotransmitter that exhibits its effects through five subtypes of dopamine receptors, important members of class A G-protein coupled receptors (GPCRs). Both subtype two (D2R) and subtype three (D3R) function via inhibition of adenylyl cyclase, and modulation of these two receptors has clinical applications in treating schizophrenia. However, the high degree of binding site conservation between D2R and D3R makes it difficult to generate pharmacological compounds that selectively bind to one but not the other. Today, we will examine how eticlopride, a D2R/D3R antagonist, binds to human D3R.

For the purposes of this exercise we will model a ligand / protein complex with a published structure, eticlopride bound to D3R (PDB: 3PBL), allowing us to compare our modeled poses with the native structure. For this tutorial we will use the crystal structure of DR3. Although, in reality it is most likely you will not have a published structure, and will have to create a comparative model for the protein (see the RosettaCM tutorial), but the steps in this tutorial will apply to both.

For this exercise, we will be preparing our input files in the **protein_prep/** and **ligand_prep/** folders. The modeling will be done in the **docking/** folder. The **scripts/** folder contains helpful ligand docking scripts that we will be using during this tutorial (you should never be copying files to or from this folder). All necessary files are also prepared in the **answers/** directory in case you get stuck.

1. Navigate to the ligand docking directory where you will find the **ligand_prep/**, **protein_prep/**, **docking/**, and **answers/** folders

```
cd ~/rosetta_workshop/tutorials/ligand_docking/1_vanilla_ligand_docking
```

2. Prepare a human dopamine 3 receptor structure. We will do this by obtaining the crystal structure (3PBL) and removing the excess information.

1. Change into the **protein_prep/** directory with the cd command

```
cd protein_prep
```

2. The **clean_pdb.py** script will allow you to automatically download a PDB file and clean it of information other than the desired protein coordinates. The 'A' option tells the script to obtain chain A only. The full crystal structure consists of two monomers.

```
~/rosetta_workshop/rosetta/tools/protein_tools/scripts/clean_pdb.py 3PBL A
```

3. There are two output files generated by clean_pdb.py: 3PBL_A.pdb contains a single chain of the protein structure and 3PBL_A.fasta contains the corresponding sequence. 3PBL_A.pdb is the receptor structure we will be using for docking, copy this file into the docking directory.

```
cp 3PBL_A.pdb ../docking
```

Note: This structure has a T4-lysozyme domain instead of the third cytoplasmic loop. The T4-lysozyme is a stabilizing feature to aid in crystallography. Normally, we would truncate this lysozyme segment and perform loop modeling (as discussed in the RosettaCM tutorial) to regenerate the intracellular loop. However in the interest of time, we will use the lysozyme containing structure as the eticlopride binding site is far from the intracellular domain.

3. Next, we will prepare the ligand files. Most of these files are already prepared for you in the interest of time, but the steps are explained.

1. cd into the directory named **ligand_prep/**

```
cd ../ligand_prep
```

2. In the directory, you will find a pair of already prepared files: eticlopride.sdf and eticlopride_conformers.sdf

1. eticlopride.sdf: This contains the eticlopride structure found in the 3PBL protein complex.

Note: You can also find the ligand file from this particular PDB structure by going to the [3PBL page](#) and scrolling down to the "Small Molecules" section. From there, you can click "Download SDF File" under the ETQ identifier.

2. `eticlopride_conformers.sdf`: This is a set of conformations for eticlopride generated outside of Rosetta. The downloaded ligand `eticlopride.sdf` file contains only the single conformation found in the PDB so we must expand the library to properly sample the conformational space. We also need to add hydrogen atoms to the model since they are not resolved in the crystal structure. Feel free to open the file in Pymol and use the arrow keys in the bottom right of the window to scroll through the different conformations:

```
pymol eticlopride_conformers.sdf
```

This particular conformational library was generated using the Meiler lab's BioChemicalLibrary (BCL). The BCL is a suite of tools for protein modeling, small molecule calculations, and machine learning. If you're interested in licensing the BCL, please visit <http://www.meilerlab.org/bclcommons> or ask one of the instructors. Other methods of ligand conformer generation include OpenEye's MOE software and web-servers such as [Frog 2.1](#) or [DG-AMMOS](#). The generated libraries will differ depending on the chosen method.

3. Generate a `.params` file and associated PDB files for eticlopride. The `.params` file contains important information about eticlopride in order to properly build the molecule in Rosetta, including the partial charges of the atoms, the connectivity of the molecule, and internal coordinates. The parameters file is necessary for ligand docking because Rosetta does not have internal records for custom small molecules in its database.

1. Type

```
~/rosetta_workshop/rosetta/main/source/scripts/python/public/molfile_to_params.py -h
```

to learn more about the script for generating the params file.

2. Type

```
~/rosetta_workshop/rosetta/main/source/scripts/python/public/molfile_to_params.py \
-n ETQ -p ETQ --conformers-in-one-file eticlopride_conformers.sdf
```

Note: You may encounter a warning about the number of atoms in the residue. This is okay as Rosetta is merely telling you that the ligand has more atoms than an amino acid.

Three total files will be generated: `ETQ.params` contains the necessary information for Rosetta to process the ligand, `ETQ.pdb` contains the first conformation, and `ETQ_conformers.pdb` contains the rest of the conformational library. I would highly suggest walking through one of these params files while looking at the corresponding PDB structure in a graphics program (Pymol, Chimera, MOE, your favorite).

4. If you use the `tail` command on `ETQ.params`, you will notice the `PDB_ROTAMERS` property line that tells Rosetta where to find the conformational library. Make sure this line has `ETQ_conformers.pdb` as the property.

```
tail ETQ.params
```

5. Now that we have the necessary files for ligand docking, let's copy them over to the docking directory.

```
cp ETQ* ../docking
```

4. Now we want to make our final preparations in the docking directory.

1. Change to the **docking/** directory

```
cd ../docking
```

2. Concatenate the ligand and protein pdb files together into one pdb file

```
cat 3PBL_A.pdb ETQ.pdb > 3PBL_A_ETQ.pdb
```

3. Open up our prepared pdb file to examine the receptor / ligand complex

```
pymol 3PBL_A_ETQ.pdb
```

4. Tip: 'all->A->preset->ligand sites->cartoon' will help you visualize the protein/ligand interface. The "Action" button is denoted by a single letter "A" in Pymol

Since this is a rudimentary exercise, we will start with the ligand in the known protein binding site. In practical application, it is unlikely that we will know the exact location of the binding site. Therefore we may need to try multiple starting locations, defining a starting point using the `StartFrom` mover or manually place the ligand into an approximate region using Pymol.

5. Once you close Pymol, make sure Rosetta has these three necessary input structure/parameter files in the docking directory. If you are missing any of these, copy them from ../answers/docking/
 1. 3PBL_A_ETQ.pdb: a single chain of the protein receptor structure with a default starting conformation for eticlopride
 2. ETQ_conformers.pdb: A pdb file containing all conformers generated from the eticlopride library
 3. ETQ.params: a Rosetta parameter file that provides the necessary properties for Rosetta to treat eticlopride
6. Next we need to make sure we have the proper RosettaScripts XML file, input options file, and “crystal complex” (the correct answer for comparison) in our directory. These files are provided to you as dock.xml, options.txt, and crystal_complex.pdb
 1. dock.xml - This is the RosettaScripts XML file that tells Rosetta the type of sampling and scoring to do. It defines the scoring function and provides parameters for both low-resolution coarse sampling and high-resolution Monte Carlo sampling.
 2. options.txt - This is the options file that tells Rosetta where to locate our input PDB structures and ligand parameters. It also directs Rosetta to the proper XML file.
 3. crystal_complex.pdb - This is the D3-eticlopride complex from the PDB. It will serve as the correct answer in our case allowing us to make comparisons between our models and actual structures.
5. Run the docking study (This should take a few minutes at most, as we’re using a reduced number of output structures):

```
~/rosetta_workshop/rosetta/main/source/bin/rosetta_scripts.linuxgccrelease \
  @options.txt -nstruct 5 -database ~/rosetta_workshop/rosetta/main/database/
```

6. The Rosetta models are saved with the prefix 3PBL_A_ETQ_ followed by a four digit identifier. Each model PDB contains the coordinates and Rosetta score corresponding to that model. In addition, the model scores are summarized in table format in the score.sc file. The two main scoring terms to consider are:
 1. total_score: the total score is reflective of the entire protein-ligand complex and is good as an overall model assessment
 2. interface_delta_X: the interface score is the difference between the bound protein-ligand complex and the unbound protein-ligand. The interface score is useful for analyzing ligand effects and for comparing different complexes.
7. One other metric to keep an eye on is the Transform_accept_ratio. This is the fraction of Monte Carlo moves that were accepted during the low resolution Transform grid search. If this number is zero or very low, the search space may be too restrictive to allow for proper sampling.
8. In benchmarking examples when we have a correct crystal structure, ligand_rms_no_super_X will give us the RMSD difference between our model ligand and the crystal structure ligand given in crystal_complex.pdb. This is an important metric when benchmarking how well your models correlate to reality. When the crystal structure is unknown, we can also calculate model RMSDs using the best scoring structure as the “true answer”.
9. Use pymol to visually compare your best-scoring model and worse-scoring model with the crystal structure provided in crystal_complex.pdb. The “all->A->preset->ligand sites->cartoon” setting in Pymol is ideal for visualizing interfaces. What interactions were successfully predicted by Rosetta?
10. The visualize_ligand.py script in the scripts directory is a useful shortcut for doing quick visualizations of protein-ligand interfaces. It takes in a PDB and generates a .pse Pymol session by applying common visualization settings. The example below shows the command lines for using this script on the 0001 model but you are free to try it on any one (or more!) of your models:

```
~/rosetta_workshop/tutorials/ligand_docking/1_vanilla_docking/scripts/visualize_ligand.py 3PBL_A_ETQ_0001.pdb
pymol 3PBL_A_ETQ_0001.pse
```

Analysis

Since we generated such a small number of structures, it is unlikely to capture all the possible binding modes that you would expect to encounter in an actual docking run. In the **docking/out/** directory, there are 500 models pre-generated using the exact same protocol. We will look at an example of how we can analyze this dataset.

1. cd into the **out/** directory

```
cd out
```

2. In addition to the 500 structures here, you will find the `score.sc`, a `score_vs_rmsd.csv` file, a `rmsds_to_best_model.data`, and several `.png` image files.

1. `score.sc`: summary score file for the 500 structures as outputted by Rosetta
2. `score_vs_rmsd.csv`: a comma separated file with the filename in the first column, `total_score` for the complex in the second column, the interface score in the third column, and ligand RMSD to the native structure in the fourth column.

This file was tabulated using the `extract_scores.bash` script and the `score.sc` file as input. This is a very specific script made for extracting useful information in ligand docking experiments. However, the script can be easily customized for extracting other information from Rosetta score files. If you have any in-depth questions about how it works or how to modify it, feel free to ask. To see how it in action, run:

```
~/rosetta_workshop/tutorials/ligand_docking/1_vanilla_docking/scripts/extract_scores.bash score.s
```

3. `rmsds_to_best_model.data`: a space separated file containing RMSD comparisons with the best scoring model (not crystal structure!) for all PDB files. A more detailed discussion of this file will come further down in the tutorial. This file has the filename in the first column, an all heavy-atom RMSD in the second column, a ligand only RMSD without superimposition in the third column, a ligand only RMSD with superimposition in the fourth column, and heavy atom RMSDs of side-chains around the ligand in the fifth column.

This file is generated using the `calculate_ligand_rmsd.py` script. It uses `pymol` to compare PDB structures containing the same residues and ligand atoms. It's a quick way of calculating the ligand RMSDs of Rosetta models. To see how this works, let's try it on the five models we generated in the previous steps:

```
cd ../
```

```
~/rosetta_workshop/tutorials/ligand_docking/1_vanilla_docking/scripts/calculate_ligand_rmsd.py \  
-n 3PBL_A_ETQ_0003.pdb -c X -a 7 -o rmsds_to_best_model.data *_000*.pdb
```

This command compares all five of your models to the one after the `-n` option. Your best scoring model may not be the one labelled 0003 so feel free to customize that option. The `-c` tells the script that the ligand is denoted as chain X. The `-a` tells the script to use 7 angstroms as the cutoff sphere for side-chain RMSDs. The `-o` option is the output file name. Lastly, we provided a list of PDBs using the wildcard selection.

The script produces the `rmsd_to_best_model.data` file that you can open in any text editor. Feel free to ask questions if you would like to discuss more of how to customize this script for your own applications. Now let's go back to the pre-generated model directory:

```
cd out
```

4. PNG files: plots made from the various data file mentioned above. Python and the `matplotlib` package was used here but you are free to use any plotting software you prefer.
3. In this case, we have the correct answer based on the crystal structure so we can examine a score vs rmsd plot to see if the better scoring models are indeed closer to the native ligand binding mode. Open up the plot with the following command:

```
gthumb score_vs_crystal_rmsd_plot.png
```

On the X-axis you will see the ligand RMSD to the ligand in the crystal structure. On the Y-axis you will see the interface delta score in Rosetta Energy Units. Notice the general correlation between RMSD and Rosetta Score, with a large cluster of highly accurate and low scoring models in the lower left hand corner.

4. In practical applications, we would not have the crystal structure for comparison. However, we can treat the best scoring model as the native model and see if we generate a similar funnel. This is one application of how we might use the `calculate_ligand_rmsd.py` script discussed earlier. Once we identify a desired “best model”, we can run the script to generate the `rmsds_to_best_model.data`. Some scripting may be required to put the information from multiple files together, depending on which software package you choose to graph with. To identify the best scoring model for this example, I selected the top 200 models based on the best overall score and then identified the best model by interface score. The best model for these plots is `3PBL_A_ETQ_0347.pdb`. Open up the first plot with:

```
gthumb score_vs_low_rmsd_plot.png
```

Again, we see a cluster of good scoring models near the best scoring model with a general downward trend further away. We can zoom in on the cluster in the lower left hand corner to get an even better picture.

```
gthumb score_vs_low_rmsd_zoom_plot.png
```

We see the same overall trend in this cluster, suggesting that the top scoring models in this run are likely to be good predictors of the true ligand binding position.

5. Finally let’s look at some structures. To sort the CSV file by interface score and take the top twenty, type:

```
sort -t, -nk3 score_vs_rmsd.csv | head -n 20
```

These should all be very low RMSD models. To compare a certain structure to the native in Pymol, use:

```
pymol 3PBL_A_ETQ_0211.pdb ../crystal_complex.pdb
```

I used `3PBL_A_ETQ_0211.pdb` as the sample structure because it is one of the best scoring models, but feel free to examine any model you like. Don’t forget the ligand site preset mode for visualizing interfaces or use the `visualize_ligand.py` script to generate pymol sessions. If you like, we can also look at some of the poor scoring models to see exactly what went wrong. To find the top 20 worse models by interface score:

```
sort -t, -nk3 score_vs_rmsd.csv | tail -n 20
```

`3PBL_A_ETQ_0424.pdb` should come up as a poor scoring, high RMSD structure. When we open it up in Pymol, we can see that the ligand binding direction is flipped 180 degrees compared to the native position. This can happen when there is an extended binding pocket, but in this case, the Rosetta score was able to discern the difference between these models.

```
pymol 3PBL_A_ETQ_0424.pdb ../crystal_complex.pdb
```

Congratulations, you have performed RosettaLigand docking study! Now use your docked models to generate hypotheses and test them in the wet lab!

Virtual High Throughput Screening (vHTS)

Conformer Generation

Conformers can be generated with a number of tools, including MOE and OMEGA. In this case, the Conformer Generation tool included as part of the BioChemical Library (BCL) suite was used. The following command can be used with the most recent version of BCL to generate conformers:

```
bcl.exe molecule:ConformerGenerator -max_iterations 1000 -top_models 100 \  
-conformation_comparer 'Dihedral(method=Max)' 30 \  
-temperature 4 -cluster -clash_weight 2.0 -clash_tolerance 0.4 \  
-ensemble_filenames 1KV2_Validation_Affinities_3D.sdf \  
-conformers_single_file 1KV2_Validation_Affinities_3D_conf.sdf
```

You can use any conformer generation tool you have available to you for this step. Your generated conformers should be output to a single SDF file. Every conformer must have 3D coordinates and hydrogens added. Conformers of the same ligand should have the same name in the SDF file. For convenience, an example conformer file is provided at `rosetta_inputs/ligands/all_ligands.sdf`.

Params File Generation

Params files contain the parameterization information for a ligand. Every ligand or Residue in a protein structure input into Rosetta must have a corresponding params file. Rosetta is distributed with a script called `molfile_to_params.py` which generates these files. However, this script is generally cumbersome for the generation of more than a small handful of ligands. The following a protocol for generating params files for large numbers of ligands:

All the scripts needed for this process are in the tools directory in the Rosetta distribution. each of the scripts below would normally be preceded by `Rosetta/tools/hts_tools`, but this directory prefix has been omitted for brevity.

1. *Split ligand files*

The conformers for all ligands are initially stored in a single SDF file, but `molfile_to_params.py` expects 1 SDF file per ligand. `sdf_split_organize.py` accomplishes this task. It takes as input a single sdf file, and will split that file into multiple files, each file containing all the conformers for one ligand. Different ligands must have different names in the sdf records, and all conformers for one ligand must have the same name. Output filenames are based on the sha1 hash of the input filename, and are placed in a directory hashed structure. Thus, a ligand with the name “Written by BCL::WriteToMDL,CHEMBL29197” will be placed in the path `./41/412d1d751ff3d83acf0734a2c870faaa77c28c6c.mol`.

The script will also output a list file in the following format:

```
ligand_id,filename  
string,string  
ligand_1,path/to/ligand1  
ligand_2,path/to/ligand2
```

The list file is a mapping of protein names to sdf file paths.

Many filesystems perform poorly if large numbers of files are stored in the same directory. The hashed directory structure is a method for splitting the generated ligand files across 256 roughly evenly sized subdirectories, improving filesystem performance.

The script is run as follows:

```
sdf_split_organize.py 1KV2_Validation_Affinities_3D_conf.sdf split_conformers/ ligand_names.csv
```

Be sure the `split_conformers/` directory exists before running the script.

2. Create Project Database

The ligand preparation pipeline uses an sqlite3 database for organization during the pipeline. The database keeps track of ligand metadata and the locations of ligand files. The project database is created using the following command:

```
setup_screening_project.py ligand_names.csv ligand_db.db3
```

An example of the project database is in `example_outputs/ligand_prep`

3. Append binding information to project database

The next step is to create a binding data file. The binding data file should be in the following format:

```
ligand_id,tag,value
string,string,float
ligand_1,foo,1.5
ligand_2,bar,-3.7
```

The columns are defined as follows:

- `ligand_id` – `ligand_id` is the name of the ligand, which must match the `ligand_id` in the `file_list.csv` file created by `sdf_split_organize.py`.
- `tag` – The name of the protein the ligand should be docked into. If a ligand should be docked into multiple proteins, it should have multiple entries in the binding data file. Note that this protocol makes a distinction between protein name, and file name. If you have 4 protein files: `foo_0001.pdb`, `foo_0002.pdb`, `bar_0001.pdb`, and `bar_0002.pdb`, then you have two proteins with the names `foo` and `bar`. The scripts expect that the protein PDB files begin with the protein name.
- `value` – The activity of the ligand. If you are doing a benchmarking study and know the activity of your ligand, you should enter it here. If you are not doing a benchmarking study, or if ligand activity is not relevant to your study, value can be set to 1.0 (or anything else). This field is currently only used in a few specific Rosetta protocols that are in the experimental stages, and is typically ignored, so it is safe to set arbitrarily in almost every case.

An example input file is provided. you can insert it into the project database with the following command:

```
add_activity_tags_to_database.py ligand_db.db3 ligand_activities.csv
```

4. Generate Params Files

The next step is to generate params files. `make_params.py` is a script which wraps around `molfile_to_params.py` and generates params files in an automated fashion. Params files will be given random names that do not conflict with existing Rosetta residue names (no ligands will be named ALA, for example). This script routinely results in warnings from `molfile_to_params.py`, these warnings are not cause for concern. Occasionally, `molfile_to_params.py` is unable to properly process an sdf file, if this happens, the ligand will be skipped. In order to run `make_params.py` you need to specify the path to a copy of `molfile_to_params.py`, as well as the path to the Rosetta database.

`make_params.py` should be run like this:

```
make_params.py -j 2 --database Rosetta/main/database \
  --path_to_params Rosetta/main/source/src/python/apps/public/molfile_to_params.py \
  ligand_db.db3 params/
```

In the command line above, the `-j` option indicates the number of CPU cores which should be used when generating params files. If you are using a multiple core machine, setting `-j` equal to the number of available cpu cores. Be sure that the `params/` directory exists before running the script.

The script will create a directory `params/` containing all params files, pdb files and conformer files.

An example of the output `params/` directory is found in `example_outputs/ligand_prep`

5. Create job files

Because of the memory usage limitations of Rosetta, it is necessary to split the screen up into multiple jobs. The optimal size of each job will depend on the following factors:

- The amount of memory available per CPU
- The number of CPUs being used
- The number of atoms in each ligand
- The number of conformers of each ligand
- The number of protein residues involved in the binding site.

Because of the number of factors that affect RosettaLigand memory usage, it is usually necessary to determine the optimal job size manually. Jobs should be small enough to fit into available memory.

To make this process easier, the `make_evenly_grouped_jobs.py` script will attempt to group your protein-ligand docking problem into a set of jobs that are sized as evenly possible. The script uses BioPython is run like this:

```
python2.7 make_evenly_grouped_jobs.py params/ protein_files/ \  
--n_chunks 1 --max_per_job 1000 --inactive_cross_dock job
```

If the script was run as written above, it would use param files from the directory `params/`, and structure files from the directory `protein_files/`. It would attempt to split the available protein-ligand docking jobs into 10 evenly grouped job files (`-n_chunks`). The script will attempt to keep all the docking jobs involving one protein system in one job file. However, if the number of jobs in a group exceeds 1000, the jobs involving that protein system will be split across multiple files (`--max_per_job`). The script will output the 10 job files with the given prefix, so in the command above, you would get files with names like “`output_prefix_01.js`”. The script will output to the screen the total number of jobs in each file. All the numbers should be relatively similar. If a job file at the beginning of the list is much larger than the others, it is a sign that you should reduce the value passed to `--max_per_job`. If the sizes of all jobs are larger than you want, increase `-n_chunks`.

Additionally, the script will take the default ligand positions from the ligand pdb files, and the protein files from the `protein_files` directory, and designate these as the “native” pose of the protein-ligand complex. This feature will allow Rosetta to compute ligand RMSDs automatically, and was used in the benchmarking studies described in the manuscript.

An example job file produced using this script is found in `example_outputs/ligand_prep`

Docking

After following the procedure above to prepare your ligands, you are ready to dock the ligands. The screening job file produced in the previous step contains the paths to the input proteins and ligands and the paths to the necessary params files. In this example, the ligand pdb's are already positioned in the ligand binding site.

The Rosetta ligand docking command should be run as follows:

```
rosetta_scripts.default.linuxgccrelease @ flags.txt \  
-in:file:screening_job_file job_01.js -parser:protocol dock.xml \  
-out:file:silent results.out
```

`flags.txt` contains flags that are always the same regardless of the input file.

This command will dock every protein-ligand binding pair and place the output in the specified silent file. In the benchmarking case described in the manual, 2000 models were made for each protein-ligand binding pair. However, in a practical application 200 models would be appropriate.

Analysis

Practical analysis

If this protocol is being used for an application project in which the correct ligand binding position is not known, the lowest scoring model for each protein-ligand binding pair should be selected. From that point, we recommend filtering by protein-ligand interface score (interface_delta_X), as well as the packstat score (Sheffler 2009) which can be computed through the InterfaceAnalyzer mover. The cutoffs for these filtering steps should depend on the range of scores present, and the number of compounds it is possible to test.

RosettaLigandEnsemble

The following protocol and commentary was directly extracted from Fu D., Meiler J. RosettaLigandEnsemble: A Small-Molecule Ensemble-Driven Docking Approach (2018).

Structure-activity relationships (SARs) refer to differences in binding affinity or biological efficacy following chemical scaffold derivatizations. Medicinal chemistry makes use of such minor modifications to optimize lead compounds for desired affinity and other pharmacological properties. This creates a massive wealth of SAR data on related ligands for a single protein target. The PubChem database alone contains over 200 million measurements of biological activities on approximately 10000 protein targets.(3) BindingDB specifically organizes a portion of its database into collections of congeneric ligands with at least one co-crystallized with the common protein target.(4) It is generally expected that highly similar ligands form similar interactions when binding to the same target.(5) We hypothesize that a docking algorithm that leverages this information can eliminate a portion of false-positive binding poses, i.e., poses that score well, but are incorrect. We have extended RosettaLigand to RosettaLigandEnsemble (RLE), an algorithm that can identify a binding mode favorable to a superimposed ensemble of congeneric ligands. This allows users to simultaneously dock a series of ligands in unison instead of individually as single ligands. We hypothesize that this will increase the efficiency and accuracy of sampling. We illustrate the hypothesized sampling advantage of RLE in Figure 1. Due to the presence of functional groups of varying sizes found within a SAR series, there may be binding modes available to certain molecules, but not others. RLE is capable of eliminating binding orientations not available to the ensemble as a whole. Furthermore, highly similar ligands are expected to bind in a similar fashion with common interactions to the chemical core

The receptor structure used in this example is the p53 core domain bound to a stabilizing small molecule (PDB: 4AGQ, protein.pdb) The ligand series should share a core scaffold by which the ligands can be aligned. This example contains fivecongeneric ligands, but any number between three and eight is a reasonable use case.

Ligand Preparation:

For clarity, these scripts will be written for a single ligand, PDBID=4AGL, however in practice, this process must be done for each ligand in the set.

PyMol pair fitting is an easy way to manually align ligands by minimizing distance between core scaffold atoms. Automated ligand alignment tools mayalso be used but generally do not perform as well compared to manual inspection. Examples of aligned ligands can be found in the /prep/aligned_ligands/ directory.

We need to again generate conformers for the ligands. Below is an template BCL script used to generate conformer files, but you can use others if you choose. Examples of generated conformers files are in prep/conformers.

```
bcl.exe molecule:ConformerGenerator -ensemble_filenames prep/aligned_ligands/4AGL.withH.sdf \
  -conformers_single_file prep/conformers/4AGL.conformers.sdf
```

Params File Preparation:

Rosetta requires params file to properly handle small molecule ligands. Prior to this step, concatenate the aligned ligand structure with the corresponding conformers into a single SDF file such that the aligned structure is first in the file. This will insure that the inputs will maintain the core scaffold alignment when generating the conformers (`cat prep/aligned_ligands/4AGL.withH.sdf prep/conformers/4AGL.conformers.sdf >> prep/make_params/make_params.4AGL.sdf`).

Since PDB files use three digit residue codes and single digit chain designations, it's helpful to assign a code for each ligand file. The example uses the prep/ligands.list file to label each ligand as residues 00B through 00F and corresponding chains B through F. This file also contains pK values for each ligand binding to the target receptor.

```
~/rosetta_workshop/rosetta/main/source/scripts/python/public/molfile_to_params.py \  
  prep/make_params/make_params.4AGL.sdf --chain B -n 00B -p 00B --conformers-in-one-file
```

For each ligand, command (3) generates a PDB file containing the single aligned ligand structure, a PDB file containing the remaining ligand conformers, and a params file containing connectivity and charge information for the ligand. Examples of these files can be found in the /prep/rosetta_inputs/ directory using the previously discussed letter designations. If you wish to incorporate SAR during docking, then use a text editor to add `NUMERIC_PROPERTY AFFINITY #####` to the end of the params file, where ##### represents an SAR measurement of the user's choice. Rank correlation is used in SAR mode and hence units only need to be self-consistent.

Input File Organization:

For RLE runs, it is convenient to prepare a single PDB file containing the aligned ligands but concatenating the individual ligand PDB files. The conformer and params do not need to be joined. This is done as ligands.pdb in the /inputs/ folder. You'll also find the previously prepared protein receptor PDB in the same directory.

The example dock.xml provided uses the settings from the benchmark. Actual application use may require the user to alter these values according to biological context. The defined scoring function is based on the existing RosettaLigand scoring function, but may be substituted in the XML script. The provided options file defines Rosetta input and output directories along with a number of sampling parameters. A full options list is available on the documentation website. The `ligand_ensemble` option is necessary to use RLE; a weight of 0 can be used to run RLE without taking SAR data into consideration.

Each simulation will produce X models, where X is the number of input ligands. These example output models are in the /outputs/ directory along with a score.sc scorefile.

```
~/rosetta_workshop/rosetta/main/source/bin/rosetta_scripts.linuxgccrelease \  
  @inputs/options -docking:ligand:ligand_ensemble 0 -nstruct 1
```

Output and Analysis:

Individual protein-ligand predicted structures are labeled by a chain and a number designation, B_1.pdb through F_1.pdb. Structures with the same numeric label are based on the same docking simulation and have a common binding pose. The protein interface contacting each ligand are optimized independently. The score.sc file contains all score terms for each simulation across a single row. Generally, individual ligand interface scores are used to rank models, with a negative score indicating a better model. These ligand interface scores are listed as `interface_delta_*`, where * is the single letter ligand chain ID. The values are appended at the end of each output PDB, and also in the scorefile for each protein-ligand pair. One suggestion is for the end user to examine the top ten percent of models for each pair.