User Guide for version 1.1

How to cite?

- Filipovic, J., Vavra, O., Plhak, J., Bednar, D., Marques, S.M., Brezovsky, J., Matyska, L., Damborsky, J., CaverDock: A Novel Method for the Fast Analysis of Ligand Transport. IEEE/ACM Transactions on Computational Biology and Bioinformatics, 2019. `https://doi.org/10.1109/TCBB.2019.2907492`

- Vavra, O., Filipovic, J., Plhak, J., Bednar, D., Marques, S.M., Brezovsky, J., Stourac, J., Matyska, L., Damborsky, J., CaverDock: A Molecular Docking-based Tool to Analyse Ligand Transport through Protein Tunnels and Channels. Bioinformatics, 2019. `https://doi.org/10.1093/bioinformatics/btz386`

Release Notes

CaverDock 1.1

- significantly improved heuristic for upper-bound computation (CaverDock generates smoother energetic profiles), improved objective function

- multiple upper-bound trajectories (differing at the beginning or the end of the tunnel) are computed and saved

- script for automatic adding of flexible residues

- pattern restraint can be applied to selected ligand atoms only (usable for docking with partially set position)

- user can control how similar is the initial snapshot to the dock_like conformation

- CaverDock reports its execution settings (for better reproducibility of experiments)

- improved parallel scaling of upper-bound computation

CaverDock 1.01

- bugfix in upper-bound trajectory computation (the bug has sometimes caused permature end of backtracking and hence worser upper-bound trajectories)

- CaverDock finishes with exit codes 0 (all OK), 1 (wrong input), 2 (cannot compute trajectories) and 3 (cannot compute upper-bound trajectory)

CaverDock 1.0

- the first version of CaverDock

1   Introduction

CaverDock is a tool designed for rapid analysis of ligand transport processes in receptors [2, 3]. It models the transportation of a ligand – a substrate, a product, an inhibitor, a co-factor or a co-solvent – from the outside environment into the receptor active or binding site (or *vice versa*).

The input for the calculation is a receptor structure in the PDB format and a ligand structure in the PDBQT format. The outputs are the ligand's trajectory and the energetic profile of the process. CaverDock implements a novel algorithm based on molecular docking that produces a contiguous ligand trajectory and estimates the binding energy along the pathway.

The current version of CaverDock uses Caver [4] for the pathway identification and a heavily modified Autodock Vina [1] as the docking engine. The tool is much faster than molecular dynamic simulations (usually 2-20 min per job), making it suitable even for virtual screening. The software is extremely easy to use since in its minimalistic configuration it requires only the setup for AutoDock Vina and the geometry of the tunnel.

2   Installation

In the current version, CaverDock is distributed as pre-compiled packages, which can be installed into the system by a standard way (e.g. using dpkg + apt). The package includes CaverDock binaries, the tunnel discretizer, scripts for a more comfortable CaverDock usage, this documentation and an example folder, containing a simple example to demonstrate the CaverDock workflow described in the Section 3.

Several tools and libraries serve as prerequisites for CaverDock, they can be installed with the following command.

Listing 1: Install pre-requisites
```
sudo apt-get install libopenmpi-dev libboost-all-dev \
python-docopt libcgal-dev python-scipy autodocktools
```

3   Quick start

The typical CaverDock workflow consists of the following steps:

- calculation of receptor's tunnels with CAVER. Selection of tunnel(s) for CaverDock analysis and their export into a PDB file;

- discretization of the tunnel(s) to a set of discs with CaverDock's tool Discretizer;

- preparation of input files for docking (by conversion of PDB to PDBQT format, optionally setting the side-chain flexibility);

- configuration and execution of CaverDock;

- visualization and interpretation of results.

In the following text, we describe the individual steps of the workflow. You may test the workflow mentioned in this section using the example input stored in the `example` folder packed with CaverDock. Any step may be omitted as the folder also contains intermediate and final results.

3.1 Tunnels calculation

The creation and selection of tunnels can be done locally by CAVER software or using the web portal [5]. For more details on CAVER usage, see CAVER User Guide [6]. You can find the already exported tunnels in the example packed with the CaverDock tool.

3.2 Tunnels discretization

Having one or more tunnels exported, they must be discretized to set of cuts (discs evenly cutting the tunnel). The discretizer tool requires a tunnel in PDB format as the input and produces a discretized tunnel for CaverDock. Please note that CAVER considers the active site as the tunnel's beginning and the surface of the receptor as the tunnel's end. This direction of the tunnel is also adopted in CaverDock. See Listing 2 for an example of the discretizer execution and Figure 1 for an illustration of discretizer output.

| Listing 2: Discretizer execution |
|---|
| discretizer.py -f tunnel1.pdb -o tunnel1.dsd |

Figure 1: Visualization of the tunnel discretization.

3.3 CaverDock execution

Finally, CaverDock must be configured and executed. The minimalistic configuration must contain a specification of the search box and names of the receptor, the ligand and the tunnel file. Please note that the search box must contain the whole ligand's trajectory (i.e. the whole tunnel) together with the area of the possible movement of side-chains. Although the configuration file can be prepared manually or with MGLTools, CaverDock package contains a script which assembles the configuration file with basic settings automatically.

Listing 3: Configuration preparation
```
cd-prepareconf -r receptor.pdbqt -l ligand.pdbqt \
-t tunnel1.dsd > caverdock.conf
```

You can include any of the CaverDock command-line options as a parameter in the configuration. For example, saving the logs from the computation to files can be turned on by calling CaverDock with command-line option `--log logName` or by adding a line `log=logName` into the configuration.

With the prepared configuration file, CaverDock may be executed. CaverDock uses MPI[1] for parallel execution. For tunnel analysis, it must be executed in at least two processes (by setting `mpirun` parameter `-np` to 2 or more).

Listing 4: CaverDock execution
```
mpirun -np 9 caverdock --config caverdock.conf --out test --log log
```

CaverDock can be run at a personal notebook or a desktop computer; its execution commonly takes from minutes to dozens of minutes. For large-scale execution or highly complex receptors and ligands, it might be more convenient to run it at a computational cluster.

CaverDock generates several output files:

- the lower-bound trajectory called `outName-lb.pdbqt`;

- the upper-bound trajectory called `outName-ub.pdbqt`;

- alternative upper-bound trajectories found by the heuristic called `outName-ub-alternatives.pdbqt`;

- the upper-bound trajectory containing the snapshot with the lowest energy called `outName-min.pdbqt`;

- a partial lower-bound or upper-bound trajectory called `outName-failed.pdbqt`, if Caver-Dock failed to find the lower-bound or upper-bound trajectory;

- a check-point file with the lower-bound trajectory called `outName-lb.txt` that can be loaded in future executions with `--load_lb outName-lb.txt`;

- a file with ranges of energies for all discs called `outName-ranges.dat`;

- a file with information about bottlenecks for each disc called `bottlenecksName.noOfDisc` if CaverDock was called with option `--dump_bottlenecks bottlenecksName`;

---

[1]Implementation of Message Passing Interface, such as OpenMPI: https://www.open-mpi.org/

- logs with detailed information about CaverDock run called `logName.noOfProcess` if CaverDock was called with option `--log logName`;

where `outName` can be set with parameter `--out outName`.

3.4 Results visualization

Files outName-lb.pdbqt and outName-ub.pdbqt are created upon successful processing of the tunnel by CaverDock. The first file contains a trajectory estimating the lower-bound of the energy profile. The lower-bound trajectory is not contiguous; however, it samples the tunnel without any gap in the ligand movement. The second trajectory is contiguous and estimates the upper-bound of the energy profile (it is the best trajectory found; however, a trajectory with better energy profile may exist). You can explore the geometry of trajectories by standard tools for viewing PDBQT files (PMV and some versions of PyMOL), and you can also create graphs of transport energy.

Energies of all snapshots as well as their positions in a tunnel are stored in PDBQT files (lines beginning with REMARK CAVERDOCK). One can easily extract the data to text with a simple script.

Listing 5: Extraction of energies

```
cd-energyprofile -d tunnel1.dsd -t outName-ub.pdbqt -s 0 > energy.dat
```

The energies should be extracted from PDBQT file containing an upper-bound trajectory (it contains the geometry of upper-bound trajectory and both upper and lower-bound energies). After the extraction, one can easily create graphs with any tool, such as Gnuplot.

Listing 6: Visualization of energies with Gnuplot

```
gnuplot -e "set xlabel \"distance\"; set ylabel \"energy\"; plot "\
"\"energy.dat\" u 1:4 w l t \"upper-bound\", \"\" u 1:6 w l "\
"t \"lower-bound\"" -p
```
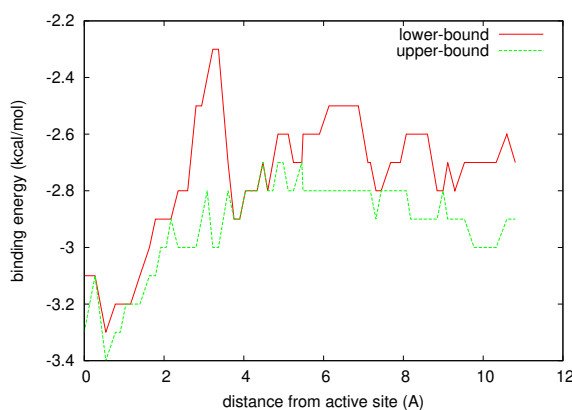

Figure 2: Visualization of energy barriers along the trajectory. The upper-bound trajectory is shown in red, the lower-bound trajectory is shown in green.

Moreover, the script cd-reportgenerator generates a LaTeXand PDF file that contains:

- details about experiment setup: CaverDock version, git commit hash, values of options

- the energy profile with lower-bound and upper-bound trajectory

- if bottleneck file is provided, the energy profile contains labels of aminoacids at the bottlenecks of the tunnel

- if the name of alternative trajectories is provided (`out-ub-alternatives.pdbqt`), the energy profile depicting profiles of all alternative trajectories

- the figure with a range of energies of viable snapshots for each disc

- experiment summary (when, where and for how long it ran)

In order to generate the report later, you need to save the log from the run by redirecting the output or by setting `--log name` and optionally save the bottlenecks if you want to examine them later in the report by setting `--dump_bottlenecks`.

The report can be generated by calling the script, creating report.tex and report.pdf.

Listing 7: Generating a report from the run

```
cd-reportgenerator -l log.0 -d tunnel1.dsd -t outName-ub.pdbqt -r outName-ranges.dat -a
    outName-ub-alternatives.pdbqt
```

4   Results interpretation

To fully understand CaverDock results, one needs to be familiar with the basic principles of CaverDock's method, which is used to analyze a ligand's movement through a tunnel. The method uses restrained docking, with two types of restraints:

- placement of a ligand atom at the 2D disc in 3D space;

- placement of the whole ligand in upper-bound vicinity of the defined snapshot.

Figure 3: A scheme depicting CaverDock method providing the lower-bound (left) and the upper-bound (right) trajectories.

First, CaverDock computes the so-called *lower-bound trajectory*. With this trajectory, a ligand atom (closest to the ligand center by default) is iteratively placed and consequent discs. This method samples the tunnel finely; however, a ligand trajectory may contain discontinuities (a ligand can, e.g. flip when moved to a different disc). Thus, some bottlenecks may not be detected (Figure 3 left). Second, CaverDock starts to search for a contiguous

trajectory by docking to consequent discs, but always in the vicinity of the previous snapshot (Figure 3 right).

Snapshots of the lower-bound trajectory are computed for each disc separately; thus, the trajectory search space is relatively small. On the other hand, the position of each snapshot of the contiguous trajectory depends on the previous snapshot. As there are exponentially many trajectories with respect to the number of discs, an exhaustive search of all contiguous trajectories is not feasible. CaverDock uses the heuristic method to prune the space of possible trajectories. The heuristic does not guarantee that the trajectory with the lowest energy is found; thus, we call a contiguous trajectory *upper-bound trajectory*, as it upper-bound energy of transport process (the actual energy is the same or lower).
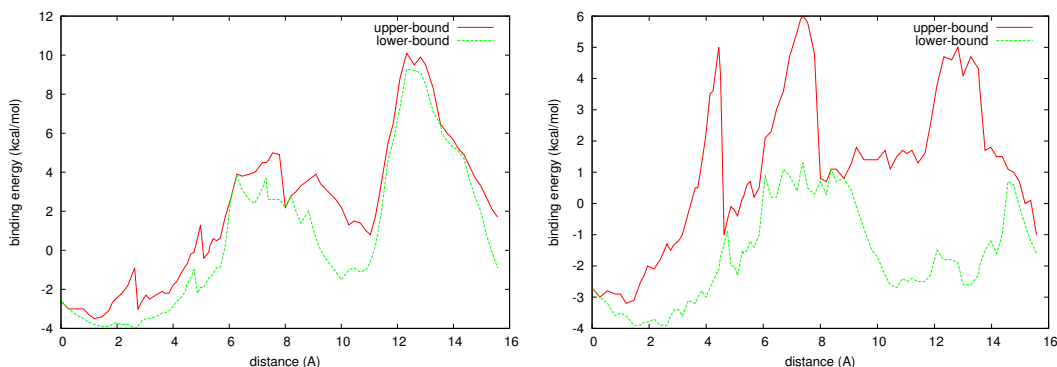

Figure 4: Examples of two distinct trajectories obtained using CaverDock. Left: similar lower and upper-bound, right: different lower and upper-bound.

Ideally, the lower-bound and upper-bound energies approximately agree (Figure 4 left). If the difference is higher than expected(Figure 4 right), we need to explore the trajectory and examine whether the bottleneck which appeared in the upper-bound trajectory is rather an artefact of the computational method, or if it is created by a real barrier in the tunnel. We can also try to improve the precision of the CaverDock computation using the procedure described in the Section 7.

5   Command line options

The main function of CaverDock is to analyze the transportation of a ligand through a tunnel. However, it can also be executed for a single docking. In the single docking mode, CaverDock performs single docking with restraints, which can be used as a building block for other tools using restrained docking.

The parameters accepted by CaverDock can be divided into three categories:

- docking options, derived from AutoDock Vina;

- options configuring tunnel analysis;

- options for restrained docking.

5.1 Docking options

CaverDock accepts the same parameters as Autodock Vina, which are used to configure, how the docking is performed. When only those parameters are given, and only one process is executed[2], CaverDock works exactly like AutoDock Vina. When a tunnel analysis or a single restrained docking is executed, docking options described in this section must also be defined, as they parametrize docking internally used in CaverDock. The list of basic docking parameters can be found in Table 1. For more information about AutoDock Vina docking and all docking parameters, we refer to AutoDock Vina Users Guide [8].

`--receptor arg`	Rigid part of the receptor (in PDBQT format).
`--flex arg`	Flexible side chains, if any (in PDBQT format).
`--ligand arg`	Ligand (in PDBQT format).

`--center_x arg`
`--center_y arg`
`--center_z arg`
	Coordinates of the center of the search box. The search box must contain space of the whole tunnel and all flexible residues with their potential movement. The CaverDock package contains a script `prepareconf.py` which computes the search box automatically (Section 3.3).

`--size_x arg`
`--size_y arg`
`--size_z arg`
	Size of the search box in Å. The search box must contain space of the whole tunnel and all flexible residues with their potential movement. The CaverDock package contains a script `prepareconf.py` which computes the search box automatically (Section 3.3).

`--out arg`	Output file(s) prefix. When docked to the tunnel, a file with the suffix `-lb.pdbqt` is generated for the lower-bound trajectory, `-ub.pdbqt` for the upper-bound trajectory and occasionally `-failed.pdbqt` if the upper-bound trajectory is not computed (the failed trajectory contains the longest known contiguous trajectory).
`--exhaustiveness arg`	Exhaustiveness of the global optimization algorithm (a higher number means a higher number of tested conformations). Usually values from the interval 1–8 are sufficient.
`--cpu arg`	Number of threads used per process.
`--log`	Writes a log file(s), usable for debugging.

Table 1: Docking parameters.

5.2 Tunnel analysis options

In this section, we introduce all parameters related to the analysis of ligand transport through a tunnel. For the transport analysis, at least two processes must be executed (using `mpirun -np 2` or greater). Note that the only required parameter is `--tunnel`, other parameters can be used to change the default settings. The basic parameters allow defining the type of the computed trajectory and the optimization strategy or hinting the initial position of a ligand (Table 2), whereas advanced parameters affect restraints applied to the docking (Table 3).

---

[2]by direct execution of the binary or executing it via `mpirun -np 1`

`--tunnel arg`	The file containing discretized tunnel. This parameter is required.
`--final_state arg`	The final state of trajectory search. May be set to LB (compute the lower-bound trajectory only) or SMOOTHED (default), computing the lower and the upper-bound trajectories.
`--optimization_strategy arg`	Strategy for the upper-bound trajectory optimization. CaverDock may minimize the highest energy in a trajectory as a primary criterion and the cumulative energetic barrier as a secondary criterion (MAX, default value) or the cumulative energetic barrier across a whole trajectory (HILL).
`--dock_like arg`	PDBQT file defining the ideal position of the ligand in the active site. The starting disc is selected to bring a ligand closest to the active site. This parameter may help when the tunnel obtained from CAVER is too deep (and thus CaverDock pushes the ligand against the energetic barrier at the bottom of the tunnel).
`--dump_bottlenecks arg`	Dump the bottlenecks observed along the trajectory. There are multiple files generated: in the file **arg**, a list of bottleneck residues and a list of bottlenecks for final upper-bound trajectory are given. For n disks, files **arg.0..**, **arg.n-1** are generated: they contain the bottlenecks for the snapshots used in the final trajectory as well as the snapshots generated during the trajectory search and not used in the output trajectory. Reporting the bottlenecks in unused snapshots may help the user to recognize which residues prevent CaverDock from moving the ligand in alternative pathways.

Table 2: Basic tunnel analysis parameters.

Whereas the basic tunnel analysis options can be set by beginner users and should not affect the robustness of the CaverDock computation, a suboptimal setting of advanced options may lead to very a slow computation or a failure of the search for contiguous conformation. Thus, these should be set carefully.

`--backtrack_threshold arg`	The energy increment over the lower-bound trajectory, which leads to the execution of backtracking. A lower value, lower energy of the upper-bound trajectory and higher computation time can be expected. The default value is 1 kcal/mol.

`--backtrack_limit arg`	The minimal number of steps (visited discs) which must be performed after backtracking, before new backtracking is executed. Values higher than 1 forbid frequent executions of backtracking from unfeasible areas. The default value is 5.
`--cont_threshold arg`	The maximal atom movement (in Angstroms) between adjacent snapshots, which is considered to be contiguous. The default value is 0.8 Å.
`--pattern_limit arg`	The maximal distance (in Angstroms) between adjacent snapshots, which is not penalized by a contiguity restraint. The default value is 0.6 Å, must be smaller than `--cont_threshold` and higher than the maximal distance of discs created during a tunnel discretization.
`--dock_like_limit arg`	The maximal distance (in Angstroms) between the first snapshot and the ideal position of the ligand (defined by `--dock_like`), which is not penalized a restraint. The default value is 2.0 Å
`--allow_flex_discontinuity`	Allows the receptor side-chain residues to perform non-contiguous movements even when the upper-bound trajectory is computed. The default behavior is that a contiguous movement is required for both the ligand and the side-chain residues.
`--parallel_workers_lb arg`	The number of processes solving each docking task of a lower-bound trajectory computation in parallel (alternative to exhaustiveness with better scaling). The default value is 4.
`--parallel_workers_smooth arg`	The number of processes solving each docking task of an upper-bound trajectory computation in parallel (alternative to exhaustiveness, scales better). The default value is 4.
`--early_exit`	Allows CaverDock to immediately halt when lower-bound at some disc cannot be computed. The parameter can be used with screening, where rapid discarding of narrow tunnels is required.

Table 3: Advanced tunnel analysis parameters.

5.3 Restrained docking options

In this section, parameters applicable to perform a single docking with restraints are introduced. Note that those parameters are not employed in the standard usage scenario of CaverDock (an analysis of the whole trajectory). However, they may be used for debugging or building a tool on the top of CaverDock.

There are essentially two types of restraints implemented: a disc, which fixes a ligand's atom and a pattern, which restricts the positions of all ligand's atoms. When the disc is

defined, docking is performed in the way that a selected ligand atom must be placed in close vicinity of the disc. The discs may define a cut of a tunnel, or a point in space (when the disc radius is set to a very small number). The pattern holds the whole ligand in the vicinity of pattern atoms. So, it is possible to search only for the ligand conformations which are close to the one defined by the pattern.

`--ccenter_x arg`
`--ccenter_y arg`
`--ccenter_z arg`
	Coordinates of the center of a disc attracting a ligand (its central atom, or the atom selected by `catomnum` parameter).

`--cnormal_x arg`
`--cnormal_y arg`
`--cnormal_z arg`
	The normal vector of a disc attracting a ligand (its central atom, or the atom selected by `catomnum` parameter).

`--cradius arg`	The radius of a disc attracting a ligand (its central atom, or atom selected by `catomnum` parameter).
`--catomnum arg`	The ID of an atom in PDB, which is attracted to a disc.
`--ctemplate arg`	PDB or PDBID of the template, which restricts the ligand atoms movement (they must remain in a vicinity of the template atoms, defined by `--ctemplate_limit`).
`--ctemplate_limit arg`	The distance (in Angstroms) of the docked ligand atoms to their counterparts in a template (defined by `--ctemplate`), which is tolerated without applying an attractive force to the template.
`--ctemplate_restrict arg`	The subset of template's atoms, which are used to restrict position of ligand atoms movement. Defined as a list of atom IDs delimited by "`_`", e.g. `1_3_4` applies the template to the first, the third and the fourth atoms only.

Table 4: Restrained docking parameters.

6 Automatic Flexibility

CaverDock 1.1 contains the experimental support for automatic addition of flexible residues. The script `cd-flexibilize` iterativelly adds flexible residdues by performing following protocol:

- compute trajectory with CaverDock;

- if convergence criterion met, finish the protocol;

- search for side-chain residues forming bottlenecks;

- construct receptor file with more flexible side-chain residues;

- repeat the protocol.

The script begins with a rigid receptor and in each iteration adds a number of flexible residues defined by `--addflex` parameter (1 by default). The number of iterations is upperbound by parameter `--maxiters` (4 by default), or the script may finish when the highest energy of the computed trajectory is under threshold defined by parameter `--stopenergy`.

Together with the configuration of flexible residues addition described above, the script requires caverdock config file passed by `-c` (must not contain receptor file), the geometry of a rigid receptor passed by `-r` and number of processes used for MPI defined by `-p`. The example of flexibilize script execution is given in Listing 8.

Listing 8: Flexibilize script execution

```
cd-flexibilize -c caverdock.conf -r receptor.pdbqt -p 3 --maxiters=2 --stopenergy=-3
```

The resulting trajectory produced by the flexibilize script is stored in files `out-lb.pdbqt` and `out-ub.pdbqt`, where `out` is the output name set by user in configuration file. The intermediate trajectories (with less flexible residues than the final trajectory) are also saved with suffix `rigid` or `flex-i` (where `i` is iteration number).

7   Best practices

7.1   Improving energy

CaverDock has been parametrized to bring a good mix of computational efficiency and precision of computation. However, there is still some room for hand-tuning the CaverDock computation. One can tune the CaverDock parameters to increase the precision of its computation and the structure of molecules to ease the computation of the ligand movement.

7.2   Improving lower-bound energy

The lower-bound energy depends on a tunnel geometry and the docking ability to find good local minima. To improve (decrease) the energy of the lower-bound trajectory, several arguments can be tuned:

- `--exhaustiveness` can be set to a higher value which increases the number of random walks of Markov-chain Monte Carlo global search algorithm and increases the probability of finding good local minima;

- `--parallel_workers_lb` can be set to a higher value and should have a similar effect to `--exhaustiveness`.

If the energy of the lower-bound trajectory is too high even with high exhaustiveness, the ligand probably cannot pass through the tunnel in the real-world, or there is some issue with the tunnel or the receptor geometry. Three typical issues are described below.

- The tunnel obtained from CAVER forces the ligand to move too deeply into the tunnel, where it reaches the energetic barrier at the tunnel bottom. This situation can be detected by exploring the energy graph: the energy is very high in the area of the beginning of the tunnel (i.e. around the position zero at x-axis). In such case, `--dock_like` parameter can be used to set the correct starting disc, so the ligand will not be pushed against the tunnel bottom.

- Side-chain residues are forming a bottleneck. We can detect this issue by exploring the bottleneck dump. In such a case, side-chain residues forming the bottleneck should be set to be flexible by MGTools [7].

- The backbone residues are forming a bottleneck. We can detect this issue by exploring the bottleneck dump. In such a case, we need to use a different geometry of the receptor (e.g. one obtained from an MD simulation). This is a typical issue when the receptor structure is taken from a crystal with tunnels closed due to intramolecular interactions and crystal packing.

7.3 Improving upper-bound energy

When we are satisfied with the lower-bound energy, we can focus on the upper-bound energy. If the difference between the lower-bound and the upper-bound energies is too high, we can tune several CaverDock parameters:

- `--backtrack_threshold` can be set to a lower value, if we consider 1 kcal/mol already as an undesired difference between the lower-bound and the upper-bound energies. Beware that tuning this value may result in longer computation times.

- `--backtrack_limit` can be set to a lower value so that CaverDock will execute backtracking more aggressively. Beware that tuning this value may result in longer computation times.

- `--multiple_search` can be set to a larger value so that CaverDock will execute more instances of forward movements or backtrackings. Beware that tuning this value may result in longer computation times if insufficient cores are used by CaverDock.

Together with the CaverDock parameters, we should also check the geometry of the tunnels and the receptor (Section 7.2), as some bottlenecks, as well as issues with a tunnel bottom, may arise only when a contiguous trajectory is computed.

7.4 Improving computation time

The CaverDock is usually quite fast on a standard desktop computer (its execution commonly takes from minutes to dozens of minutes). However, the execution time can be improved by the following actions.

- Use as low flexible side-chain residues as possible. The computational time of a docking grows rapidly with the number of degrees of freedom of the system. Using flexible side-chain residues may greatly improve the energy profile; however, we recommend setting flexibility only on residues, which are proved to form a bottleneck in a particular tunnel (i.e. they are reported with `--dump_bottlenecks` option).

- Use the right number of parallel processes. CaverDock uses MPI; thus, it can use multiple cores of a desktop machine or even multiple nodes of a cluster. The number of processes (passed by *-np* parameter of `mpirun`) should be set to a number of virtual cores plus one (e.g. use `-np 9` on a machine with 8 cores). Please note that computing lower-bound trajectory scales very well (CaverDock can utilize a hundred of cores in a typical scenario), whereas upper-bound trajectory scales up to the number of concurrently executed docking multiplied by concurrent searches (set by `--parallel_workers_smooth` and `--multiple_search`, respectively).

- Try increasing the number of parallel workers instead of exhaustiveness. The default number of processes solving the same docking scenario in parallel is 4. If one wants to increase the exhaustiveness of docking, the values of `--parallel_workers_lb` and `--parallel_workers_smooth` parameters may be increased instead of the value of `--exhaustiveness`.

The tips described above improve the CaverDock speed in general. However, some issues result from the properties of analyzed biochemical systems. We summarize typical issues and suggest possible solutions in the listing below.

- Backtracking is executed very frequently (this can be observed in CaverDock log files). When an upper-bound trajectory is computed, CaverDock tries to keep its energy as close to the lower-bound as possible. In some cases, it is not possible, and CaverDock executes a lot of backtracking trying to find a better trajectory without any success. The number of executed backtracking runs can be decreased by parameters `--backtrack_threshold` (higher difference of the lower-bound and upper-bound energies may be required to start backtracking) or `--backtrack_limit` (the frequency of backtracking execution may be lowered). Beware that usually high computational times and a high number of backtracking executions are related to suboptimal geometry of the tunnel or the receptor (Section 7.2 for geometry optimization tips).

- The number of degrees of freedom is very high. If the number of flexible side-chain residues is high, we can fix some residues in such a position which does not form a bottleneck and make them rigid. We may also try to compute a contiguous ligand movement with a non-contiguous movement of side-chain residues by using parameter `--allow_flex_discontinuity` (which removes some restraints in the search space). If the origin of the high number of degrees of freedom is mainly the complexity and the size of a ligand, this system may be too complex to be efficiently analyzed by CaverDock.

8   Troubleshooting

8.1   CaverDock cannot compute the lower-bound trajectory

There are several possible reasons, which may prevent CaverDock from computing the lower-bound trajectory.

- A part of the tunnel is too narrow and therefore it forms a strong repulsive barrier. This situation needs manual inspection: CaverDock saves only the part of the lower-bound trajectory, which was successfully computed. The missing part of the trajectory probably contains a strong repulsive barrier such as a residue preventing the ligand from moving through the tunnel. The receptor geometry needs to be fixed by adding flexibility to side-chain residues or using a different snapshot of the receptor.

- The ligand is forced to move against the active site bottom. This problem arises from a geometrical analysis of the tunnel where the geometrical approximation of the tunnel is too deep. An easy solution is to use the `--dock_like` parameter to navigate CaverDock where to start with the tunnel analysis.

- The ligand is too complex to be successfully docked. CaverDock can be re-executed with higher exhaustiveness or a higher number of parallel workers.

- The computation fails due to the stochastic nature of CaverDock. In this case, starting CaverDock once again should solve the problem.

8.2   CaverDock cannot compute the upper-bound trajectory

When an upper-bound trajectory cannot be computed, the user should inspect the lower-bound trajectory first. When the lower-bound trajectory already contains a high energetic barrier, the tunnel or the receptor geometry needs to be modified (Section 7.2). If it is not possible to improve the lower-bound trajectory, the selected ligand is likely not able to pass through the tunnel.

When the lower-bound trajectory does not contain any significant barrier, and the upper-bound trajectory is still not computed, CaverDock may not able to analyze the trajectory because of a high ligand complexity, or there is a bottleneck not detected by the lower bound trajectory. This type of bottlenecks can be found by inspecting the lower-bound trajectory in the vicinity of the disc, where the computation of the upper-bound trajectory has failed[3]. There should be visible non-contiguities in the lower-bound trajectory, which overpass the bottleneck (e.g. ligand flip, as shown in Fifure 3). The bottleneck needs to be manually identified and fixed (e.g. by using flexible side-chain residues).

9   FAQ

9.1   Is CaverDock execution deterministic?

In the default settings, it is not deterministic (i.e., re-executing CaverDock may result in different trajectories). There is, however, a way how to make CaverDock deterministic. It must be executed in two processes only (using `mpirun -np 2`), and the random number generator must be set to a constant seed (using `--seed x`, where x is any number, which must be the same for all deterministic executions). Please note that when CaverDock is executed only for single docking (possibly with restraints), setting a constant seed is sufficient.

9.2   Why does CaverDock return two types of energy?

A heuristics drives the CaverDock computation of an upper-bound trajectory. It means that CaverDock cannot guarantee that the contiguous trajectory is optimal. Thus, CaverDock is computing also a lower-bound trajectory: a scenario, which can be unrealistically optimistic due to non-continuities in the trajectory. The real energy (with respect to the given force-field and the geometry of the input molecules) is at most as high as the energy of the upper-bound trajectory and at least as low as the energy of the lower-bound trajectory.

References

[1] O. Trott, A. J. Olson, AutoDock Vina: Improving the Speed and Accuracy of Docking with a New Scoring Function, Efficient Optimization and Multithreading. *Journal of Computational Chemistry 31 455-461 (2010).*

---

[3]CaverDock reports a number of the last disc, for which the upper-bound trajectory was computed. The number of snapshots in the lower-bound trajectory is equal to the number of the disc.

[2] Vavra, O., Filipovic, J., Plhak, J., Bednar, D., Marques, S.M., Brezovsky, J., Stourac, J., Matyska, L., Damborsky, J., CaverDock: A Molecular Docking-based Tool to Analyse Ligand Transport through Protein Tunnels and Channels. Bioinformatics, 2019. https://doi.org/10.1093/bioinformatics/btz386

[3] Filipovic, J., Vavra, O., Plhak, J., Bednar, D., Marques, S.M., Brezovsky, J., Matyska, L., Damborsky, J., CaverDock: A Novel Method for the Fast Analysis of Ligand Transport. IEEE/ACM Transactions on Computational Biology and Bioinformatics, 2019. https://doi.org/10.1109/TCBB.2019.2907492

[4] E. Chovancova, A. Pavelka, P. Benes, O. Strnad, J. Brezovsky, B. Kozlikova, A. Gora, V. Sustr, M. Klvana, P. Medek, L. Biedermannova, J. Sochor, J. Damborsky, CAVER 3.0: A Tool for Analysis of Transport Pathways in Dynamic Protein Structures. *PLOS Computational Biology 8: e1002708 (2012).*

[5] CAVER Web Portal. https://loschmidt.chemi.muni.cz/caverweb/

[6] CAVER User Guide. http://www.caver.cz/fil/download/manual/caver_userguide.pdf

[7] MGLTools. http://mgltools.scripps.edu/

[8] AutoDock Vina Users Guide. http://vina.scripps.edu/manual.html