DOCKLAB · SELF-STUDY COURSE

Molecular Docking & Virtual Screening

Lecture notes and study material

N–H···O 2.9 Å

12 lectures · worked examples · self-tests with answers

Real results from tested runs of the DockLab labs:

Abl kinase + imatinib (1IEP) · HIV-2 protease + indinavir (1HSG) · EGFR + erlotinib (1M17)

Tools: AutoDock Vina 1.2.7 · AutoDock4 · Meeko 0.8 · Open Babel 3.2 · RDKit · PyMOL · PLIP

Edition 1 · November 2026 · companion to the DockLab course page

Contents

Contents 2

Before you start 3

Drug discovery and where docking fits 4

Protein structures and the Protein Data Bank 6

Receptor preparation 8

Ligand preparation 10

Search space and grid boxes 12

Search algorithms and scoring functions 14

Validation: RMSD and redocking 17

Analysing poses and interactions 19

Virtual screening and compound libraries 22

Ranking, enrichment and pitfalls 24

Publication figures and reporting 28

Capstone project guide 30

Command reference 31

Glossary 32

Answers to the self-tests 33

Software and references 35

HOW TO USE THESE NOTES

Before you start

These notes are the reading material for the DockLab four-week course. Each lecture matches one module on the course page and is followed by a hands-on lab. Read a lecture once quickly, then a second time with a notebook, writing the key ideas in your own words. Then run the lab.

Every lecture has the same parts:

- **Learning objectives**: what you should be able to do afterwards.
- **Lecture text**, with figures, equations and tables.
- **Worked examples** (amber boxes): calculations and decisions done step by step.
- **Key points** (teal boxes): the ideas to remember.
- **Self-test**: three or four questions. Answers are in Appendix C.

Real data, not invented numbers

Every result quoted in these notes (docking scores, RMSD values, run times, enrichment factors, interaction lists) comes from running the DockLab labs on an Apple Silicon Mac with the exact commands from the course page. Your numbers may differ slightly with other software versions or random seeds; the trends should not.

Study schedule

Week	Lectures	Labs	Practice
1 · 2–7 Nov	1 Drug discovery · 2 Structures · 3 Receptor prep · 4 Ligand prep · 5 Search space	0, 1, 2 (prep)	Sheet A
2 · 9–14 Nov	6 Search & scoring · 7 Validation · 8 Interactions	2 (run), 3	Sheet B
3 · 16–21 Nov	9 Virtual screening	4, 5, 6, 7 (prep)	Sheet C
4 · 23–28 Nov	10 Enrichment · 11 Figures · 12 Capstone	7, 8, capstone	Sheet D

UNITS AND CONVENTIONS

Energies are in kcal/mol (1 kcal/mol = 4.184 kJ/mol). Distances are in ångström (1 Å = 0.1 nm). More negative docking scores mean more favourable predicted binding. RT = 0.593 kcal/mol at 298 K.

LECTURE 1

Drug discovery and where docking fits

LEARNING OBJECTIVES

- Describe the stages of small-molecule drug discovery and where computational methods are used.
- Distinguish structure-based from ligand-based design.
- State what a docking program does and what its output means.
- List the main approximations and limitations of docking.

1.1 The drug discovery pipeline

A new small-molecule medicine typically takes 10–15 years from idea to approval, and most candidates fail along the way. The work moves through recognisable stages (Figure 1.1). A **target**, usually a protein whose activity drives a disease, is identified and validated. **Hits**, molecules with measurable activity against the target, are found by high-throughput screening, fragment screening or *virtual screening*. Hits are improved into **leads** and then optimised for potency, selectivity and drug-like properties before preclinical and clinical testing.

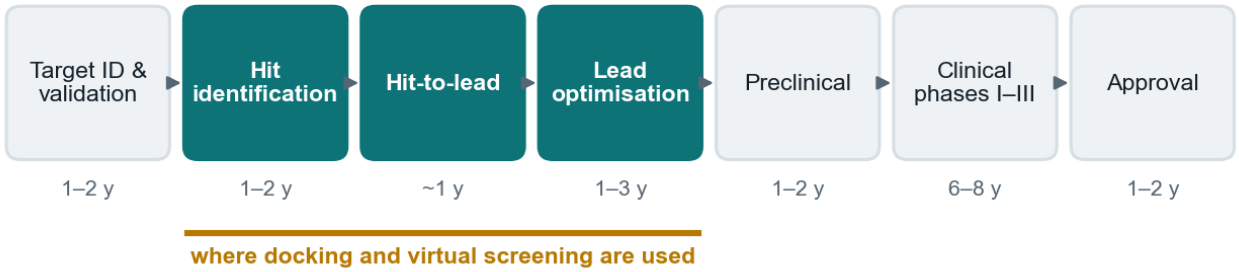


Figure 1.1 · The small-molecule pipeline with approximate durations. Docking is used mainly in hit identification and the optimisation stages (teal).

Computer-aided drug design (CADD) is used throughout, but docking earns its place early: it can evaluate millions of molecules for the cost of a computer cluster, and it gives a structural hypothesis (a binding pose) that chemists can reason about.

1.2 Structure-based and ligand-based design

Structure-based drug design (SBDD) uses the three-dimensional structure of the target, from X-ray crystallography, cryo-EM, NMR or prediction. Docking, structure-based pharmacophores and free-energy calculations belong here. **Ligand-based drug design (LBDD)** uses only known active molecules: similarity searching, QSAR models and ligand pharmacophores. The two are complementary; when a good structure exists, docking is usually the first tool used.

1.3 What docking does

Molecular docking predicts how a small molecule (the **ligand**) sits in a binding site of a protein (the **receptor**) and estimates how strongly it binds. It answers two questions:

- **Pose:** where and in which orientation and conformation does the ligand bind?
- **Score:** how favourable is that pose, expressed as an estimated binding free energy (kcal/mol)?

Every docking program combines a **search algorithm**, which generates candidate poses by moving, rotating and twisting the ligand, with a **scoring function**, which ranks them. Lecture 6 covers both in

detail. The workflow you will use in every lab follows the same five steps (Figure 1.2).



Figure 1.2 · The docking protocol used throughout the course.

1.4 Models of binding

Fischer's **lock-and-key** model treats protein and ligand as rigid complementary shapes. Koshland's **induced fit** lets the protein change shape on binding, and **conformational selection** says the protein already samples many shapes and the ligand picks one. Standard docking is closest to lock-and-key for the protein (rigid receptor) and fully flexible for the ligand's rotatable bonds. That compromise is the root of many docking failures.

1.5 Limitations you must remember

- **Rigid receptor.** Side chains and loops are frozen. A ligand that needs the pocket to open will be missed.
- **Approximate scoring.** Scores correlate only weakly with measured affinity (typical errors of 2 kcal/mol or more). They are useful for ranking and enrichment, not for predicting an exact IC_{50}.
- **Crude solvent and entropy.** Water is mostly implicit; the entropy lost on binding is estimated from the number of rotatable bonds.
- **Fixed chemistry.** Protonation states and tautomers are whatever you supplied; covalent bonds and metal coordination need special treatment.

> **KEY POINTS**
>
> Docking = search + scoring. Output = poses + scores.
>
> Use docking to generate hypotheses and to enrich a library in actives, not to prove that a molecule works.
>
> A docking result is only as good as the structure and preparation behind it.

Self-test

1. Name the two outputs of a docking calculation and the two components of every docking program.
2. At which pipeline stages is docking most useful, and why?
3. Give two reasons why a molecule with an excellent docking score might be inactive in the laboratory.

LECTURE 2

Protein structures and the Protein Data Bank

LEARNING OBJECTIVES

- Judge the quality of an experimental structure from its resolution, R-free and B-factors.
- Read the records of a PDB file.
- Choose a suitable structure for docking.

2.1 Where structures come from

Method	What to check	Notes for docking
X-ray crystallography	Resolution (≤ 2.5 Å good), R-free (< 0.25), B-factors, electron density for the ligand	Most docking structures. Hydrogens are not seen; protonation must be added.
Cryo-EM	Global and local resolution in the pocket	Large complexes and membrane proteins; side chains uncertain below ~3.5 Å.
NMR	Ensemble of models; restraint violations	Several models: can be used for ensemble docking.
Prediction (AlphaFold, homology)	pLDDT in the pocket; template identity	Usually apo-like: pockets may be closed or side chains misplaced. Validate before use.

Resolution measures how fine the detail in the experimental data is: at 1.5 Å individual atoms are distinct, at 3 Å side-chain positions become guesses. **R-free** measures how well the model predicts data it was not fitted to. The **B-factor** (temperature factor) of each atom reflects how much it moves or how poorly it is defined; high B-factors in the binding site are a warning sign.

2.2 Anatomy of a PDB file

A PDB file is fixed-width text. Header records describe the experiment; coordinate records list the atoms. Protein atoms are **ATOM** records; everything else (ligands, waters, ions, buffer molecules) is **HETATM**. A real line from the Abl kinase structure used in Lab 2:

```
         1         2         3         4         5         6         7
12345678901234567890123456789012345678901234567890123456789012345678901234567890
ATOM      1  N   MET A 225      19.353  41.547  -3.887  1.00 72.26           N
```

Columns	Field	Example above
1–6	Record type	ATOM
7–11	Atom serial number	1
13–16	Atom name	N (backbone nitrogen)
17	Alternate location (altLoc)	blank (only one position)
18–20	Residue name	MET
22	Chain	A
23–26	Residue number	225
31–54	x, y, z coordinates (Å)	19.353, 41.547, −3.887
55–60	Occupancy	1.00
61–66	B-factor (Å²)	72.26 (high: a flexible N-terminus)

Other records to know: **REMARK 465** lists missing residues and **REMARK 470** missing atoms; **SEQRES** gives the full sequence; **CONECT** lists extra bonds. Remember that the file is the *asymmetric unit* of the crystal. It may contain several copies of the protein (1IEP has chains A and B) or only half of a biological dimer (HIV protease is a dimer of identical chains). The *biological assembly* files on RCSB give the functional form.

2.3 Choosing a structure for docking

1. Correct protein and species; no mutations in or near the binding site.
2. A **holo** structure (with a bound ligand) in the site you want, ideally with a ligand similar to yours.
3. Best resolution available, complete side chains in the pocket, low B-factors there.
4. The right functional state: for kinases, active (DFG-in) or inactive (DFG-out) conformations bind different inhibitors.

WORKED EXAMPLE · the three structures used in this course

PDB	Protein · ligand	Method · resolution	Notes
1IEP	Abl kinase · imatinib (STI)	X-ray · 2.10 Å	Two copies (A, B). Imatinib binds the inactive DFG-out state.
1HSG	HIV-2 protease · indinavir (MK1)	X-ray · 2.0 Å	Homodimer (chains A + B form one site). Conserved flap water. HIV-2, not HIV-1.
1M17	EGFR kinase · erlotinib (AQ4)	X-ray · 2.60 Å	One chain. Active-like conformation.

Self-test

1. A structure has 3.4 Å resolution and B-factors above 90 Å² in the pocket. Should you dock into it? What would you look for instead?
2. Which PDB columns hold the residue name, chain and residue number?
3. Why is 1HSG a single binding site even though it contains two chains?

LECTURE 3

Receptor preparation

LEARNING OBJECTIVES

- Decide which waters, ions and cofactors to keep.
- Explain protonation states and use pK_a values to choose them.
- Describe the PDBQT format and AutoDock atom types.

3.1 The preparation checklist

1. Keep the chain(s) that form the binding site; remove other copies.
2. Remove the co-crystallised ligand (save it separately: it defines the site and the validation reference).
3. Remove bulk waters and crystallisation additives (glycerol, sulfate, PEG, DMSO).
4. Keep structural waters, metal ions and cofactors that are part of the site.
5. Repair missing side-chain atoms; resolve alternate locations (keep the higher-occupancy one).
6. Add hydrogens at the chosen pH, checking ionisable residues in the site.
7. Assign charges and atom types, and write the receptor as PDBQT.

3.2 Waters

Most crystallographic waters are bulk solvent and should be removed, because a fixed water occupies space a ligand might use. A few are **structural**: buried, conserved across many structures, and bridging protein and ligand through hydrogen bonds. The classic example is the HIV protease *flap water*, which links the backbone NH groups of Ile50 and Ile50′ to the inhibitor. Lab 6 tests the effect of keeping it.

3.3 Hydrogens and protonation

X-ray structures contain no hydrogens, yet hydrogens decide which atoms are H-bond donors. Ionisable residues must be set to the state that dominates at the working pH. The fraction of a group that is protonated follows the Henderson–Hasselbalch equation:

$$\text{fraction protonated} = 1 / (1 + 10^{\text{pH} - \text{p}K_a})$$

Residue	Typical pK_a	State at pH 7.4	Watch for
Asp / Glu	3.9 / 4.3	Deprotonated (–)	Catalytic pairs (aspartyl proteases) where one is protonated
His	6.0	Neutral, mostly	Three forms: HID (Nδ1–H), HIE (Nε2–H), HIP (both, +). Flip by H-bond pattern.
Lys / Arg	10.5 / 12.5	Protonated (+)	Buried lysines can be neutral
Cys	8.3	Neutral	Catalytic cysteines (thiolate) and disulfides

Tools such as **PROPKA** predict pK_a values from the structure; PDB2PQR, PDBFixer and H++ apply them when adding hydrogens.

WORKED EXAMPLE · the HIV-2 protease catalytic aspartates (Lab 6 result)

PROPKA 3 predicted pK_a = 5.07 for Asp25 in chain A and 9.06 for Asp25 in chain B.

At pH 7.4, chain A: $1 / (1 + 10^{7.4 - 5.07}) = 1 / (1 + 214) \approx 0.5\%$ protonated (charged).

Chain B: $1 / (1 + 10^{7.4 - 9.06}) = 1 / (1 + 0.022) \approx 98\%$ protonated (neutral).

Conclusion: the two aspartates share one proton, as the enzyme mechanism requires. A tool that simply deprotonates every Asp at pH 7.4 gets this site wrong.

3.4 Charges, atom types and PDBQT

AutoDock programs read **PDBQT** files: PDB columns plus a partial charge (Q) and an AutoDock atom type (T). Non-polar hydrogens are merged into their carbons; polar hydrogens (type HD) are kept because they make hydrogen bonds.

Type	Meaning	Type	Meaning
C	aliphatic carbon	A	aromatic carbon
N	nitrogen, not an acceptor (e.g. amide N–H)	NA	nitrogen H-bond acceptor
OA	oxygen acceptor	SA	sulfur acceptor
HD	polar hydrogen (donor)	Zn, Mg, Fe…	metal ions

AutoDock4 uses the charges (Gasteiger) in its electrostatic term. **Vina ignores charges** but needs the atom types. In our test, the Open Babel route (`obabel -xr`) wrote all receptor charges as zero: harmless for Vina, wrong for AutoDock4. That is why Lab 3 prepares its receptor in AutoDockTools, which adds Gasteiger charges.

KEY POINTS

Remove what does not belong; keep what the ligand interacts with.

Protonation decides H-bond donors and acceptors: check catalytic and buried residues by hand.

Vina: types matter, charges don't. AutoDock4: both matter.

Self-test

1. List three things you remove and two things you might keep when preparing a receptor.
2. A histidine Nε2 points at a ligand carbonyl oxygen 2.9 Å away. Which tautomer should it be?
3. A Glu in a hydrophobic pocket has a predicted pK_a of 7.9. What fraction is protonated at pH 7.4?
4. Why does Vina not need partial charges?

LECTURE 4

Ligand preparation

LEARNING OBJECTIVES

- Obtain ligands from public databases and convert between formats.
- Generate sensible 3D structures with correct protonation, tautomers and stereochemistry.
- Explain the torsion tree and why rotatable bonds make docking harder.

4.1 Where ligands come from

Database	Content	Typical use
PubChem	> 100 million compounds with 3D conformers for many	Single compounds by name or CID; bioassay data
ChEMBL	Bioactive molecules with measured potencies	Known actives for a target; SAR
ZINC	Purchasable compounds, ready-protonated, in tranches	Virtual screening libraries
DrugBank	Approved and investigational drugs	Drug repurposing screens
COCONUT	Natural products	Natural-product screens
DUD-E / DUDE-Z	Actives with property-matched decoys	Benchmarking (Lab 7)

4.2 Representations

A **SMILES** string describes a molecule's graph in one line: aspirin is `CC(=O)Oc1ccccc1C(=O)O` (lower-case letters are aromatic atoms, digits close rings, parentheses are branches, @ and @@ give stereochemistry). SMILES has no coordinates. **SDF/MOL** files hold 3D coordinates with bond orders and can contain many molecules separated by `$$$$`. **MOL2** adds atom types and charges.

4.3 From 2D to a docking-ready 3D structure

1. **Generate 3D coordinates.** Open Babel's `--gen3d` or RDKit's ETKDG build a reasonable conformer from SMILES.
2. **Set protonation and tautomers** for pH 7.4. Amines become charged, carboxylic acids deprotonated. Check the result by eye.
3. **Check stereochemistry.** A wrong stereocentre is a different molecule. If it is undefined, enumerate the isomers.
4. **Minimise** with a force field (MMFF94, UFF) to remove strained bonds and clashes. Docking keeps bond lengths and angles fixed, so any distortion stays in every pose.
5. **Build the torsion tree and write PDBQT** with Meeko or AutoDockTools.

WORKED EXAMPLE · checking an automatic protonation (from Lab 2)

Open Babel's `-p 7.4` produced imatinib with **both** piperazine nitrogens protonated (the SMILES in the PDBQT header shows two [N@H+]). In reality, only the N-methyl nitrogen ($pK_a \approx 8$) is charged at pH 7.4; the benzylic nitrogen is much less basic.

Vina gives charges little weight, so redocking still succeeded (0.25 Å). For AutoDock4, or for interpreting salt bridges, you would fix this with a pK_a-aware tool (e.g. Dimorphite-DL, OpenEye, ChemAxon) or by editing the SMILES.

4.4 The torsion tree

In the ligand PDBQT, a rigid **ROOT** fragment is surrounded by **BRANCH**es, one per rotatable bond. The search algorithm changes the torsion angles of these bonds. A bond is rotatable if it is a single bond, not in a ring, not terminal (rotating a methyl changes nothing) and, by default, not an amide C–N bond (which is planar). The start of the real imatinib file from Lab 2:

```
REMARK SMILES Cc1ccc(NC(=O)c2ccc(C[N@H+]3CC[N@H+](C)CC3)cc2)cc1Nc1nccc(-c2cccnc2)n1
REMARK H PARENT 6 4 25 13 14 39 17 41
ROOT
ATOM      1  N   UNL     1      16.600  51.810  14.798  1.00  0.00    -0.322 N
ATOM      2  C   UNL     1      15.629  51.747  15.784  1.00  0.00    +0.255 C
ATOM      3  O   UNL     1      14.789  52.612  15.983  1.00  0.00    -0.269 OA
ATOM      4  H   UNL     1      17.326  51.148  14.863  1.00  0.00    +0.170 HD
ENDROOT
BRANCH   1   5
ATOM      5  C   UNL     1      16.678  52.727  13.688  1.00  0.00    +0.043 A
ATOM      6  C   UNL     1      15.822  54.641  12.322  1.00  0.00    +0.046 A
ATOM      7  C   UNL     1      16.950  54.502  11.393  1.00  0.00    -0.027 A
```

Imatinib has **7** rotatable bonds; indinavir has **14**. Each extra torsion multiplies the size of the conformational search space, so flexible ligands need more sampling (Lecture 6) and dock less reproducibly. Above about 20–30 torsions, docking becomes unreliable; peptides and macrocycles need special care (Meeko can treat macrocycle rings as flexible).

Self-test

1. Write the steps to turn a SMILES string into a docking-ready PDBQT file, naming one tool for each.
2. Which of these is rotatable: a C–C single bond in a benzene ring, a C–CH_3 bond, an amide C–N bond, or the C–N bond between a phenyl ring and a piperazine?
3. Why must ligands be minimised before docking?

LECTURE 5

Search space and grid boxes

LEARNING OBJECTIVES

- Define a docking box from a bound ligand or from site residues.
- Convert between Vina and AutoDock4 box definitions.
- Judge when blind docking is acceptable.

5.1 Focused and blind docking

The search algorithm only explores a box. In **focused docking** the box surrounds a known site, usually the position of a co-crystallised ligand plus a margin; this is what you should do whenever the site is known. In **blind docking** the box covers the whole protein. It sounds attractive but the search space grows enormously, and the scoring function can be fooled by surface grooves.

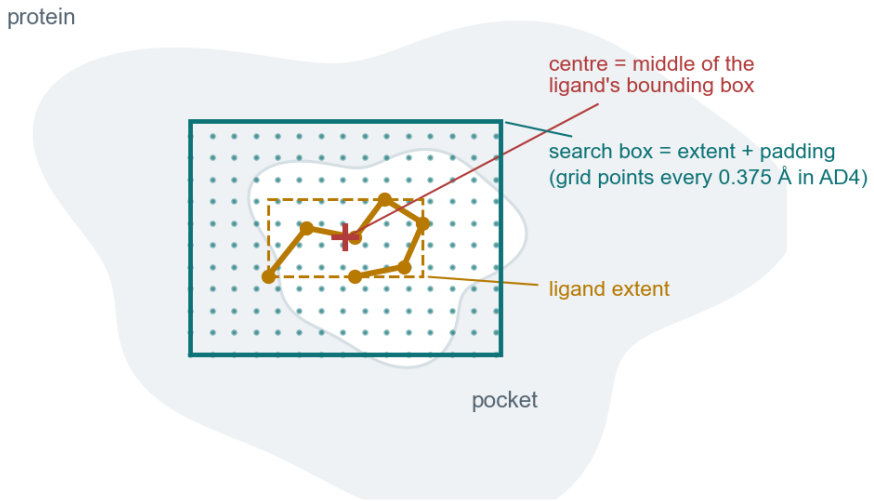


Figure 5.1 · Defining the box. The centre is the middle of the ligand's bounding box; the box size is the ligand's extent plus padding on each side.

5.2 Box definitions

	AutoDock Vina	AutoDock4 (AutoGrid)
Centre	center_x, center_y, center_z (Å)	gridcenter x y z (Å)
Size	size_x, size_y, size_z in Å	npts x y z (number of points, must be even) × spacing
Resolution	internal, not set by user	spacing, default 0.375 Å
Pre-computed maps	no (computed on the fly)	yes: one .map file per ligand atom type + electrostatics + desolvation
Warning	search volume > 27 000 Å³	very large maps use memory and disk

AutoDock4 box edge (Å) = npts × spacing e.g. 60 × 0.375 Å = 22.5 Å

WORKED EXAMPLE · the imatinib box from Lab 2

Bounding box of the crystal ligand → centre (15.190, 53.903, 16.917); extent 8.7 × 16.7 × 13.5 Å.

Adding 5 Å padding on each side (10 Å in total) gives a Vina box of 18.7 × 26.7 × 23.5 Å, volume = 18.7 × 26.7 × 23.5 ≈ 11 700 Å³ (well below 27 000 Å³).

For AutoDock4 at 0.375 Å spacing: 18.7 / 0.375 = 49.9, 26.7 / 0.375 = 71.2, 23.5 / 0.375 = 62.7 → round up to even numbers: **npts 50 72 64**.

5.3 Finding a site without a ligand

- Literature: mutagenesis data, known catalytic or allosteric residues.
- Superimpose a homologous holo structure and copy the ligand position.
- Pocket detection: fpocket, P2Rank, CASTp, DoGSiteScorer. They rank cavities by size and shape; the top pocket is usually, but not always, the real one.

5.4 Box size in practice (Lab 5 result)

Imatinib was docked with four boxes, exhaustiveness 8 and the same seed (42). Times are inflated because a screen was running at the same time.

Box	Size (Å)	Top score	Pose-1 RMSD	Best RMSD (any pose)	Time
tight	12.7 × 20.7 × 17.5	-13.336	0.24	0.24	62 s
standard	18.7 × 26.7 × 23.5	-13.273	0.25	0.25	47 s
large	30.0 × 30.0 × 30.0	-10.883	12.28	9.85	59 s
blind	57.9 × 49.5 × 69.4	-13.307	0.33	0.33	43 s

The tight and standard boxes found the crystal pose. The 30 Å box **failed** with this seed: its best pose is flipped (12.3 Å) and scores about 2.4 kcal/mol worse. The blind box, with a volume of about 199 000 Å³ (seven times Vina's 27 000 Å³ warning level), happened to succeed. Do not conclude that bigger boxes are better: a single run of a stochastic search is partly luck. Larger boxes make failure more likely, so they need higher exhaustiveness and several seeds before you trust the answer.

Self-test

1. Convert a Vina box of 24 × 24 × 24 Å into AutoDock4 npts at 0.375 Å spacing.
2. Why does a very large box need higher exhaustiveness?
3. Name two ways of placing a box when no ligand-bound structure exists.

LECTURE 6

Search algorithms and scoring functions

LEARNING OBJECTIVES

- Explain the Lamarckian genetic algorithm and Vina's Monte Carlo search.
- Write down the terms of the AutoDock4 and Vina scoring functions.
- Choose sampling settings from evidence (Lab 4).

6.1 Search algorithms

A pose is described by the ligand's position (3 numbers), orientation (3) and torsion angles (one per rotatable bond). Imatinib therefore has 6 + 7 = 13 degrees of freedom, indinavir 20. Searching such spaces exhaustively is impossible, so docking programs use stochastic methods.

Lamarckian genetic algorithm (AutoDock4)

A population of poses (default 150) evolves over generations. Each pose's 'genes' are its translation, orientation and torsions. The fittest (lowest energy) poses reproduce with crossover and mutation. A fraction undergo **local search** (Solis–Wets), and the improved pose is written back into the genes: inheritance of acquired traits, hence 'Lamarckian'. Each independent **run** (ga_run) returns one final pose; typical settings are 10–100 runs with 2 500 000 energy evaluations each (ga_num_evals). The poses are then **clustered** by RMSD (default tolerance 2.0 Å). A large, low-energy cluster suggests a well-converged answer.

Monte Carlo with local optimisation (Vina)

Vina runs several independent chains of an iterated local search: a random change to position, orientation or a torsion, followed by quasi-Newton (BFGS) optimisation using the gradient of the score; the step is accepted by the Metropolis criterion. **exhaustiveness** (default 8) is the number of independent chains, run in parallel on CPU cores. Vina returns up to **num_modes** poses (default 9) within **energy_range** (default 3 kcal/mol) of the best, which is why Lab 2 returned only 6 poses.

6.2 Scoring functions

Family	Principle	Examples
Force-field based	Physics terms (van der Waals, electrostatics) with empirical weights	AutoDock4, DOCK
Empirical	Simple interaction terms fitted to measured affinities	Vina, Vinardo, ChemScore, GlideScore
Knowledge-based	Statistical potentials from atom-pair frequencies in known complexes	DrugScore, PMF, ITScore
Machine learning	Models trained on structures and affinities	RF-Score, GNINA (CNN), many recent models

AutoDock4

The AutoDock4 free energy is a weighted sum of pairwise terms plus a torsional penalty:

$$\Delta G = W_{vdW}\cdot\Sigma(A/r^{12} - B/r^{6}) + W_{hbond}\cdot\Sigma\, E(t)(C/r^{12} - D/r^{10})$$
$$+ W_{elec}\cdot\Sigma\, q_i q_j/(\varepsilon(r)\cdot r) + W_{desolv}\cdot\Sigma(S_i V_j + S_j V_i)e^{-r^2/2\sigma^2} + W_{tor}\cdot N_{tor}$$

The terms are a 12–6 Lennard-Jones dispersion/repulsion, a directional 12–10 hydrogen-bond term (E(t) depends on angle), Coulomb electrostatics with a distance-dependent dielectric ε(r), a desolvation term

based on atomic volumes, and a penalty for conformational entropy lost when N_{tor} rotatable bonds are frozen. AutoDock4 reports the estimated ΔG and the matching K_i.

AutoDock Vina

Vina's terms depend on the **surface distance** $d = r - R_i - R_j$ between two atoms of radius R_i, R_j ($d = 0$ means the atoms just touch):

Term	Function of d	Weight	Applies to
gauss1	$\exp(-(d/0.5)^2)$	−0.0356	all heavy-atom pairs
gauss2	$\exp(-((d-3)/2)^2)$	−0.00516	all heavy-atom pairs
repulsion	d^2 if $d < 0$, else 0	+0.840	all heavy-atom pairs
hydrophobic	1 if $d < 0.5$; linear to 0 at $d = 1.5$	−0.0351	hydrophobic atom pairs
hydrogen bond	1 if $d < -0.7$; linear to 0 at $d = 0$	−0.587	donor–acceptor pairs

$$\text{score} = \Sigma\,(\text{weighted terms}) \,/\, (1 + 0.0585 \cdot N_{rot})$$

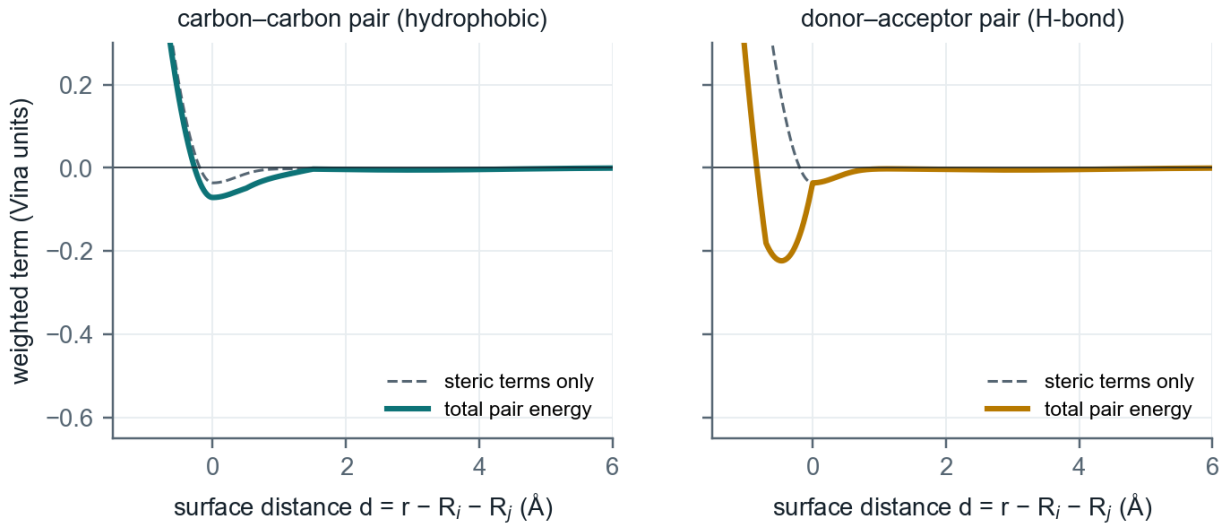


Figure 6.1 · Vina's pair energy for two atom pairs, computed from the published terms and weights. The well near d = 0 rewards close contact; the hydrogen-bond term makes donor–acceptor contacts about ten times stronger than carbon–carbon contacts.

Vina 1.2 can also use the **Vinardo** function (`--scoring vinardo`) or the AutoDock4 function (`--scoring ad4`, which needs AutoGrid maps). Comparing them is a quick robustness check.

6.3 How much sampling is enough? (Lab 4 result)

Lab 4 redocked imatinib at six exhaustiveness values with three random seeds each (Figure 6.2).

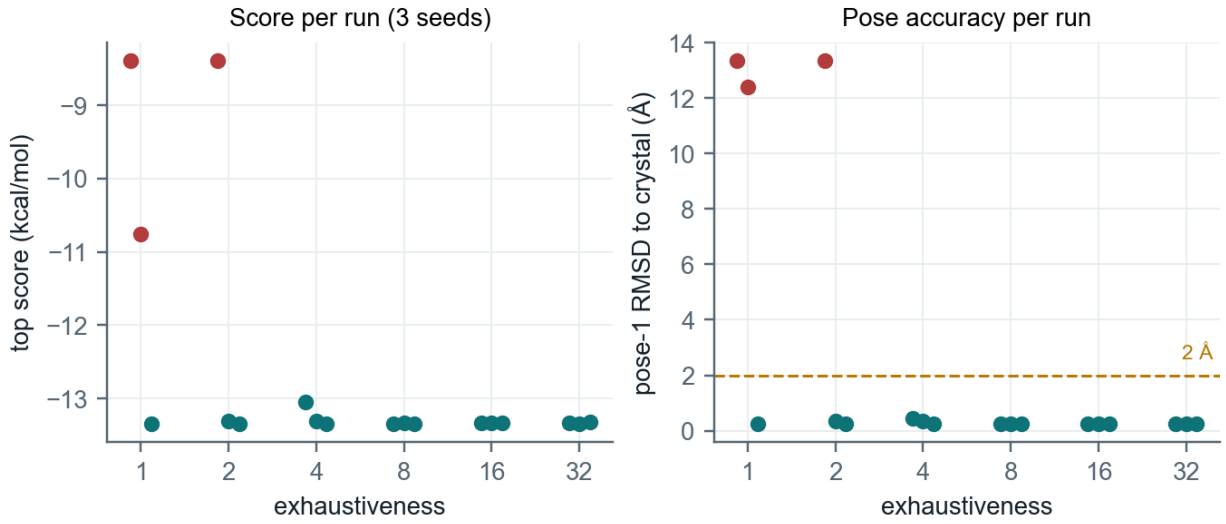


Figure 6.2 · Real Lab 4 data. Red = pose-1 RMSD above 2 Å (wrong pose). At exhaustiveness 1–2, three of six runs returned a flipped pose scoring −8.4 to −10.8 kcal/mol; from 4 upward every run found the crystal pose (−13.0 to −13.35 kcal/mol, RMSD ≤ 0.46 Å).

The lesson: with too little sampling, the search simply fails to find the best pose, and the reported score is meaningless. Once sampling is sufficient, extra effort costs time without changing the answer. For this ligand, exhaustiveness 8 (the default) is safe; flexible ligands, large boxes and final production runs justify 16–32. Always test with more than one seed.

KEY POINTS

AutoDock4: Lamarckian GA, many runs, clusters. Vina: Monte Carlo + BFGS, exhaustiveness = independent chains.

Scores are weighted sums of simple terms plus a flexibility penalty; they are not measured free energies.

Under-sampling looks like a bad score and a wrong pose. Test sampling with several seeds.

Self-test

1. How many degrees of freedom does a ligand with 9 rotatable bonds have during docking?
2. What does 'Lamarckian' mean in AutoDock4's search algorithm?
3. Two Vina runs of the same ligand with different seeds give −7.1 and −9.4 kcal/mol. What is the most likely problem, and the fix?
4. Why does Vina divide the score by $(1 + 0.0585 \cdot N_{rot})$?

LECTURE 7

Validation: RMSD and redocking

LEARNING OBJECTIVES

- Calculate and interpret RMSD between poses.
- Run and judge a redocking experiment.
- Tell sampling failures from scoring failures.

7.1 Root-mean-square deviation

RMSD measures how far the atoms of a docked pose are from the same atoms in a reference (usually the crystal pose):

$$\mathrm{RMSD} = \sqrt{\left(1/N\right) \cdot \Sigma_{i=1..N} \left|r_i^{\mathrm{pose}} - r_i^{\mathrm{ref}}\right|^2}$$

For docking, three rules apply. Use **heavy atoms only**. Compute it **in place**, without superimposing the two ligands: alignment would hide exactly the error we want to measure. Make it **symmetry-corrected**: a phenyl ring flipped 180° is the same pose, and a naive atom-by-atom RMSD would penalise it. Open Babel's `obrms` does all three; use `-f` to compare the reference with every pose.

WORKED EXAMPLE · small RMSD by hand

Three atoms are displaced by 0.5, 1.0 and 1.5 Å. RMSD = $\sqrt{((0.25 + 1.00 + 2.25)/3)} = \sqrt{1.167} = 1.08$ Å.

One large error dominates: displacements of 0.2, 0.2 and 3.0 Å give $\sqrt{((0.04 + 0.04 + 9)/3)} = 1.74$ Å.

7.2 Redocking

Redocking removes the crystal ligand, docks it back into its own receptor and measures the RMSD of the result. An RMSD ≤ 2.0 Å is the accepted threshold for success. Passing shows that your preparation, box and settings can reproduce a known answer. Do it before every new project.

Lab 2 result: imatinib redocked into Abl kinase (1IEP), exhaustiveness 8, seed 42:

Pose	Vina score (kcal/mol)	RMSD to crystal (Å)	Verdict
1	-13.27	0.25	correct
2	-11.27	12.28	wrong (flipped)
3	-10.92	13.44	wrong (flipped)
4	-10.56	13.20	wrong (flipped)
5	-10.46	12.06	wrong (flipped)
6	-10.35	13.66	wrong (flipped)

Pose 1 reproduces the crystal structure almost exactly (0.25 Å) and scores -13.27 kcal/mol, about 2 kcal/mol better than any alternative. The other poses sit 12–14 Å away: they place imatinib in the same pocket but reversed end to end. Here the scoring function clearly prefers the right answer.

7.3 Diagnosing failures

Observation	Diagnosis	What to do
No pose within 2 Å anywhere in the list	**Sampling failure** (or a preparation error)	Raise exhaustiveness, check box and protonation, check the ligand geometry

Observation	Diagnosis	What to do
A correct pose exists but is not ranked first	**Scoring failure**	Rescore with another function, consider waters or protonation, use consensus
Correct pose ranked first in some seeds only	Under-sampling	More exhaustiveness, several seeds
Pose correct but score poor	Missing interaction (water, metal, covalent bond)	Model the missing feature

7.4 Beyond redocking

Redocking is optimistic, because the receptor is already shaped for that ligand. **Cross-docking** (docking ligand X into a structure solved with ligand Y) is a harder and more realistic test. **Ensemble docking** uses several receptor structures to represent flexibility. For screening, the decisive test is **enrichment** of known actives over decoys (Lecture 10).

Lab 6 result: the effect of one water

Receptor	Seed	Score (kcal/mol)	RMSD (Å)
no water	1	-11.54	0.52
no water	2	-11.55	0.50
no water	3	-11.53	0.41
with flap water	1	-11.96	0.49
with flap water	2	-11.96	0.51
with flap water	3	-11.99	0.50

Indinavir redocks correctly (≤ 0.52 Å) with or without the flap water, so here the water is not needed to find the pose. Keeping it improves the score by about 0.4 kcal/mol, because Vina counts the extra hydrogen bonds. That is well inside scoring error. For a ligand that does not displace or use the water, keeping it could matter much more.

Self-test

1. Pose atoms are displaced by 1, 1, 1 and 3 Å. Calculate the RMSD.
2. Why must docking RMSD not be computed after superimposing the ligands?
3. Redocking gives pose 1 at 4.8 Å and pose 3 at 1.1 Å. Is this a sampling or a scoring failure? What would you try?

LECTURE 8

Analysing poses and interactions

LEARNING OBJECTIVES

- Identify the main types of protein–ligand interaction and their geometric criteria.
- Interpret an interaction profile in its biological context.
- Use ligand efficiency to compare molecules of different size.

8.1 Interaction types

Interaction	Typical geometry (approx. PLIP defaults)	Strength / notes
Hydrogen bond	donor–acceptor 2.5–3.5 Å (up to 4.1 Å), angle > 100°	Directional; the backbone NH and C=O of the kinase hinge are classic partners
Salt bridge	opposite charges ≤ 5.5 Å	Strong in buried sites, weaker at the solvent-exposed surface
Hydrophobic contact	carbon–carbon ≤ 4.0 Å	Weak individually; many contacts add up; drives much of binding
π–π stacking	ring centroids ≤ 5.5 Å, parallel or T-shaped	Aromatic rings: Phe, Tyr, Trp, His
Cation–π	cation to ring centroid ≤ 6.0 Å	Lys/Arg with aromatic ligand rings, or the reverse
Halogen bond	C–X···O/N ≤ 4.0 Å, near-linear	Cl, Br, I acting as electrophiles
Metal coordination	≤ 3.0 Å	Zn^{2}□ in metalloproteases, Mg^{2}□ in kinases
Water bridge	ligand–water–protein, each leg 2.5–4.1 Å	See the HIV protease flap water

8.2 Reading an interaction profile: imatinib in Abl kinase

PLIP analysis of the redocked imatinib pose (Lab 8) found five hydrogen bonds. Each matches a known feature of how imatinib binds:

Residue	D–A (Å)	Role in the kinase
Met318	3.01	**Hinge** backbone NH: the anchor H-bond shared by almost all ATP-site kinase inhibitors
Thr315	2.97	**Gatekeeper**. The T315I mutation removes this H-bond and causes clinical imatinib resistance
Glu286	2.95	αC-helix glutamate
Asp381	2.91	Asp of the **DFG motif**; imatinib binds the inactive 'DFG-out' state
Ile360	3.00	Backbone contact near the piperazine, which extends towards solvent

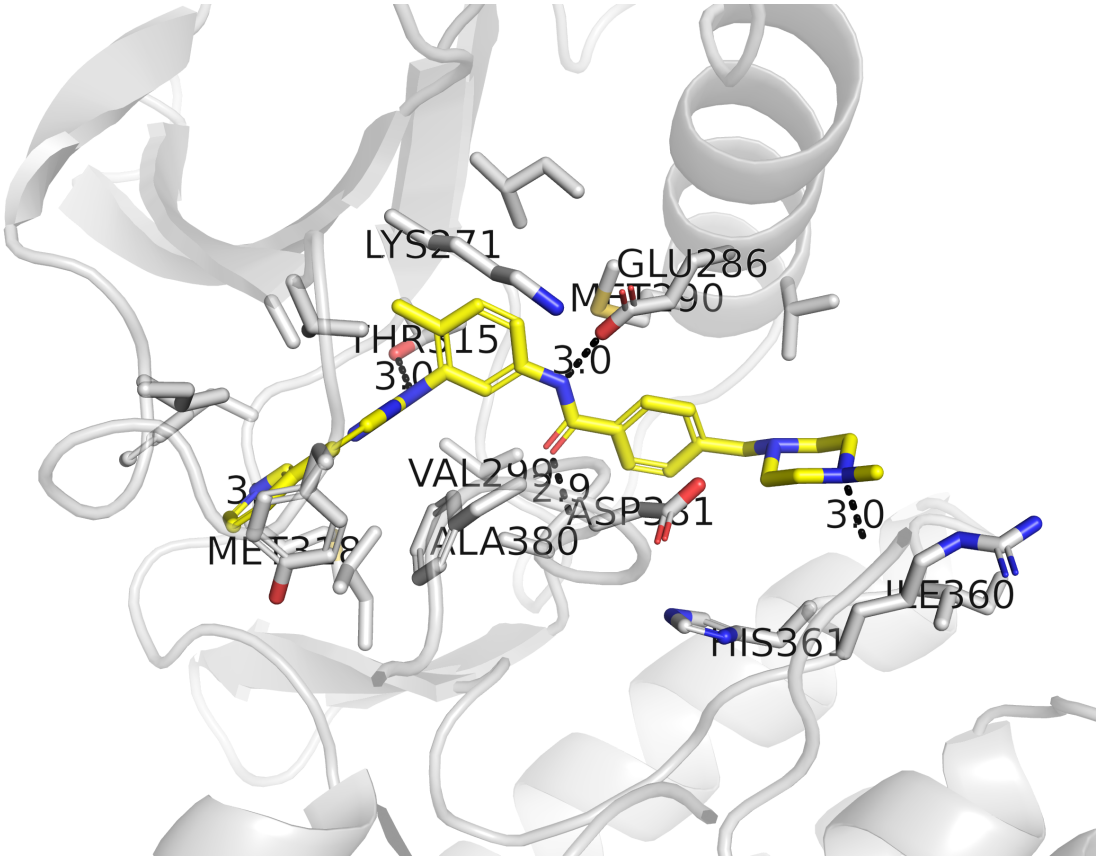


Figure 8.1 · Redocked imatinib (yellow) in Abl kinase, rendered with the Lab 8 PyMOL script. Dashed lines are polar contacts with heavy-atom distances in Å.

This is what a convincing pose looks like: it reproduces interactions that biology and medicinal chemistry already tell us matter. A pose with a good score but none of the expected interactions deserves suspicion.

8.3 Ligand efficiency

Bigger molecules make more contacts and so score better whether or not they fit well. **Ligand efficiency (LE)** divides the binding energy by the number of heavy (non-hydrogen) atoms, N_{heavy}:

$$LE = -\Delta G / N_{heavy} \text{ (kcal/mol per heavy atom; } \geq 0.3 \text{ is good)}$$

Related metrics: **LLE** = pIC_{50} − logP (potency beyond lipophilicity, ≥ 5 desirable) and **BEI** = pIC_{50} / MW (kDa).

WORKED EXAMPLE · ligand efficiency of imatinib

Vina score of the redocked pose: -13.27 kcal/mol. Imatinib ($C_{29}H_{31}N_7O$) has 29 + 7 + 1 = 37 heavy atoms.

LE = 13.27 / 37 = 0.36 kcal/mol per heavy atom, comfortably above 0.3.

A 20-heavy-atom fragment scoring −7.0 kcal/mol has LE = 0.35: just as efficient, despite a far weaker score.

8.4 Converting scores to affinities, carefully

The binding free energy relates to the dissociation constant by $\Delta G = RT \ln K_d$, so $K_d = \exp(\Delta G/RT)$. Every 1.36 kcal/mol is a factor of ten in K_d at 298 K. Figure 8.2 shows the scale. Treat converted docking scores as orders of magnitude at best.

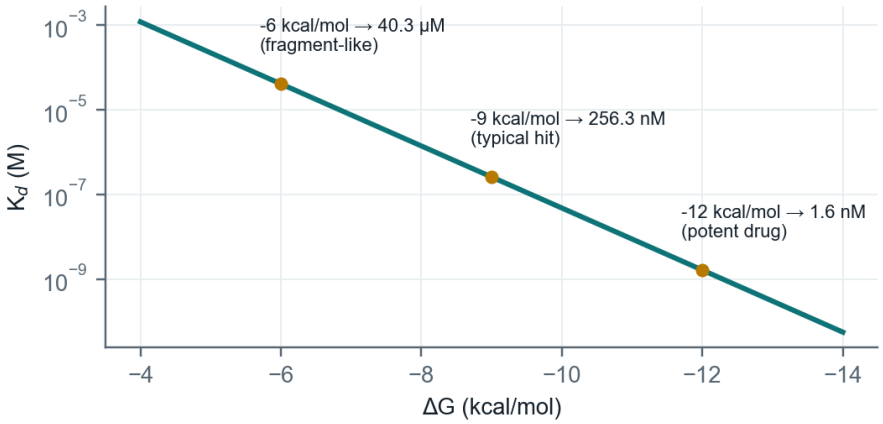


Figure 8.2 · K_d = exp(ΔG/RT) at 298 K. A 2 kcal/mol scoring error is a factor of about 30 in K_d.

8.5 Visual inspection checklist

- Does the pose make the interactions known from crystal structures and SAR (e.g. the hinge H-bond for a kinase)?
- Are polar groups buried without partners? Are charged groups paired?
- Are there clashes, or hydrophobic groups pointing into solvent?
- Is the ligand conformation reasonable (no eclipsed bonds, planar amides)?
- Do several top poses agree? Does the pose survive a different seed?

Self-test

1. List four interaction types with a typical distance criterion for each.
2. Why is the Thr315 hydrogen bond clinically important for imatinib?
3. Compound A scores −10.2 kcal/mol with 41 heavy atoms; B scores −8.1 with 25. Calculate LE for both. Which would you prioritise?
4. Convert −10.0 kcal/mol to K_d at 298 K.

LECTURE 9

Virtual screening and compound libraries

LEARNING OBJECTIVES

- Design a virtual screening workflow from library to hit list.
- Apply drug-likeness and interference filters.
- Estimate the computing time a screen needs.

9.1 The screening funnel

Virtual screening docks a library and keeps the best-ranked fraction for closer inspection and purchase. A typical funnel: library (10^4–10^9) → property filters → docking → top 1–5% → interaction filters and visual inspection → clustering for diversity → 20–100 compounds bought and tested → a handful of confirmed hits. Hit rates of a few percent are a success.

9.2 Filters

Filter	Rule	Purpose
Lipinski Rule of 5	MW ≤ 500, logP ≤ 5, H-bond donors ≤ 5, acceptors ≤ 10	Oral absorption (guideline, not law)
Veber	rotatable bonds ≤ 10, polar surface area ≤ 140 Å²	Oral bioavailability
PAINS	substructure alerts (Baell & Holloway, 2010)	Remove compounds that interfere with assays
Reactive / toxic groups	e.g. aldehydes, Michael acceptors, acyl halides	Avoid false positives and liabilities
Practical	remove salts, duplicates, metals; limit charge	Clean input for docking

9.3 Benchmark sets: actives and decoys

To test whether docking works for a target before screening a real library, mix known **actives** with **decoys**: molecules presumed inactive but matched to the actives in molecular weight, logP, charge and H-bond counts. Property matching stops a method from 'winning' by simply preferring big or charged molecules. DUD-E (102 targets) and its successor DUDE-Z (43 targets) supply such sets; DEKOIS and LIT-PCBA are alternatives. Lab 7 uses 10 approved EGFR inhibitors and 213 DUDE-Z EGFR decoys.

9.4 The library preparation pipeline

```
obabel decoys200.smi -O decoys200.sdf --gen3d                          # SMILES -> 3D
cat actives.sdf decoys200.sdf > library.sdf                               # one multi-molecule SDF
obabel library.sdf -O library_H.sdf -p 7.4                                # protonate
obabel library_H.sdf -O library_min.sdf --minimize --ff MMFF94 --steps 500   # minimise
mk_prepare_ligand.py -i library_min.sdf --multimol_outdir ligs            # one PDBQT per molecule
```

In the test run, all 223 molecules converted, with one stereo warning from Open Babel for a decoy whose stereocentres it could not correct. Read the warnings: they flag exactly the molecules that may be wrong.

9.5 Computing cost

Docking time scales with the number of ligands, their flexibility, the box size and exhaustiveness. A rough estimate:

$$\text{total time} \approx N_{\text{ligands}} \times (\text{time per ligand}) / (\text{number of parallel jobs})$$

WORKED EXAMPLE · planning a screen

In Lab 2, one Vina run for imatinib took about 10 s on an 8-core laptop (exhaustiveness 8, all cores in use).

The Lab 7 screen of 223 molecules took 24 minutes (≈ 6.5 s per ligand) on the same laptop.

1 000 compounds × 10 s ≈ 2.8 hours. 100 000 compounds ≈ 12 days on one laptop, or about 3 hours on 100 cluster nodes.

For millions of compounds, GPU programs (AutoDock-GPU, Vina-GPU, Uni-Dock) or cloud resources become necessary.

9.6 Tools

PyRx wraps Open Babel and Vina in a graphical interface and is convenient for a few hundred ligands. A **shell loop** (Lab 7) is transparent, resumable and runs anywhere. Vina 1.2 also has a built-in batch mode (`--batch` with `--dir`). On clusters, split the ligand list into chunks and submit one job per chunk.

Self-test

1. Draw the virtual screening funnel with approximate numbers at each stage.
2. Why must decoys be property-matched to the actives?
3. Estimate the time to dock 20 000 ligands at 15 s each on 4 parallel jobs.

LECTURE 10

Ranking, enrichment and pitfalls

LEARNING OBJECTIVES

- Calculate enrichment factors and interpret ROC curves.
- Recognise the common pitfalls in ranking docking results.
- Plan sensible follow-up for docking hits.

10.1 Enrichment factor

If a screen ranks N compounds containing A actives, and the top x% (n compounds) contains a actives, the enrichment factor is the hit rate in the top set divided by the hit rate of random selection:

$$EF_{x\%} = (a / n) / (A / N)$$

EF = 1 is random. The maximum possible is N/A (if A ≤ n). $EF_{1\%}$ and $EF_{5\%}$ are reported because, in a real campaign, only the top of the list is ever tested.

10.2 ROC curves and AUC

A receiver operating characteristic (ROC) curve walks down the ranked list and plots the fraction of actives found against the fraction of decoys found. The area under it (AUC) is the probability that a randomly chosen active ranks above a randomly chosen decoy: 0.5 is random, 1.0 is perfect. AUC weights the whole list equally, so early-recognition metrics such as BEDROC and logAUC are also used.

10.3 The EGFR screen (Lab 7 result)

223 ligands (10 actives, 213 decoys) were docked into EGFR (1M17) with Vina (exhaustiveness 8, seed 42).

Metric	Value	Meaning
$EF_{1\%}$	0.0	0 active(s) in the top 2 compounds
$EF_{5\%}$	0.0	0 actives in the top 11
$EF_{10\%}$	2.0	2 actives in the top 22
ROC AUC	0.73	probability that an active outranks a decoy

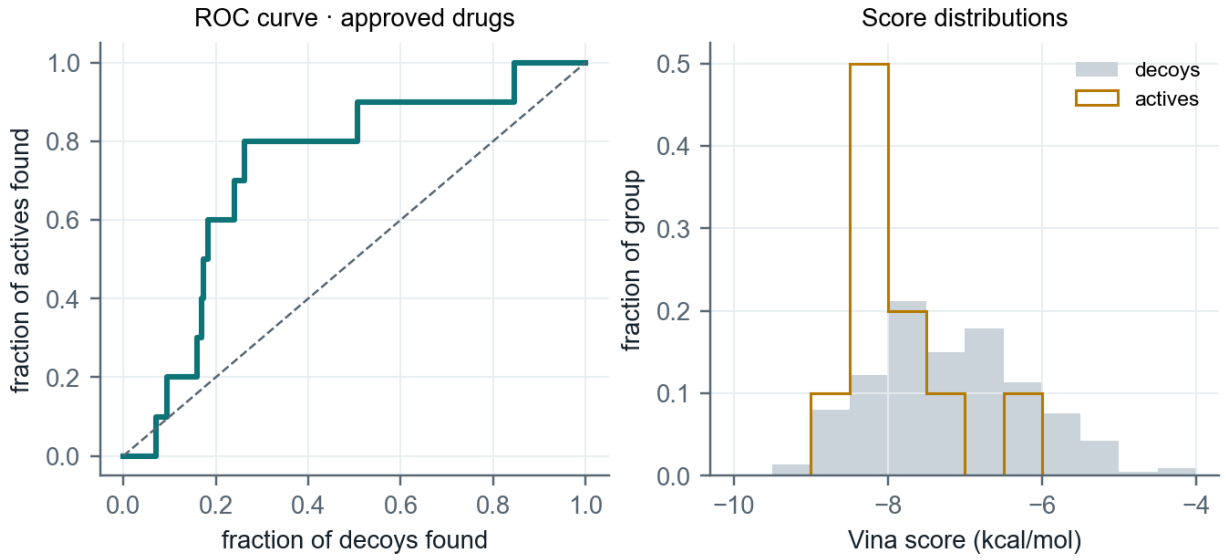


Figure 10.1 · Real Lab 7 results. Left: ROC curve (dashed = random). Right: Vina scores of actives and decoys.

Where each active ranked:

Active	Rank	Score (kcal/mol)	Binding mode in reality
dacomitinib	16 / 223	-8.68	covalent (Cys797)
lapatinib	22 / 223	-8.49	binds an inactive αC-out conformation
vandetanib	37 / 223	-8.24	reversible, ATP site
gefitinib	40 / 223	-8.18	reversible, ATP site
afatinib	42 / 223	-8.16	covalent (Cys797)
canertinib	45 / 223	-8.14	covalent (Cys797)
neratinib	58 / 223	-7.89	covalent (Cys797)
osimertinib	64 / 223	-7.88	covalent (Cys797)
erlotinib	117 / 223	-7.27	reversible, ATP site
icotinib	190 / 223	-6.08	reversible, ATP site

10.4 Case study: a screen that skipped validation

An AUC of 0.73 says actives tend to rank above decoys, but $EF_{1\%}$ and $EF_{5\%}$ of zero mean that a chemist buying the top 5% would have bought no actives at all. Most striking, erlotinib, the very ligand crystallised in 1M17, ranked 117 of 223. That should have been caught before screening. Redocking the crystal erlotinib into the same receptor gave:

Receptor preparation	Seed 42	Seed 1	Seed 2	Passes
All waters removed (as screened)	5.79 Å	5.80 Å	5.96 Å	0 / 3
Bridging water HOH 10 kept	1.18 Å	8.55 Å	1.62 Å	2 / 3

Without waters, the correct pose was found (pose 2, 1.55 Å) but ranked below a wrong one: a scoring failure. In the crystal structure, erlotinib's quinazoline N3 binds to Thr766 through a water molecule; removing it removes an interaction the scoring function needs. Keeping that single water makes the redock pass in two of three seeds. The protocol should be fixed, and validated with several seeds, *before* the screen is run. Lab 7 on the course page now includes that step.

Rerunning the whole screen with the water kept (same 223 molecules, same settings and seed) gave:

Receptor	ROC AUC	$EF_{1\%}$	$EF_{5\%}$	Erlotinib rank
All waters removed	0.73	0.0	0.0	117 / 223
Bridging water kept	0.77	0.0	0.0	104 / 223

The water fixed the pose (the screened erlotinib now docks at 1.75 Å from the crystal pose) but barely changed the ranking. The second problem is the **benchmark itself**. DUDE-Z decoys are property-matched to DUDE-Z's own EGFR actives, which are mostly small molecules, not to the ten approved drugs used here. Our actives average 34 heavy atoms; the decoys average 20. Ranking by raw score favours big molecules, ranking by ligand efficiency favours small ones (AUC 0.02 here), and neither tells us whether docking recognises EGFR inhibitors. A fair benchmark uses the actives the decoys were built for (DUDE-Z's `ligands.smi`). Testing whether known drugs rank highly is a useful, separate exercise, but it needs decoys matched to those drugs.

10.5 The fair benchmark

Finally, the screen was run the right way: all 68 unique DUDE-Z EGFR actives with every 8th DUDE-Z decoy (532), docked into the validated receptor with the bridging water. Actives and decoys now have similar size (23 and 20 heavy atoms on average). The DUDE-Z SMILES are already protonated, so the `-p 7.4` step was skipped. The 600 dockings took 66 minutes.

EGFR screen	Molecules	ROC AUC	$EF_{1\%}$	$EF_{5\%}$	$EF_{10\%}$
Approved drugs, no water	223	0.73	0.0	0.0	2.0
Approved drugs, water kept	223	0.77	0.0	0.0	1.0
Matched DUDE-Z set, water kept	600	**0.72**	**0.0**	**1.5**	**2.1**

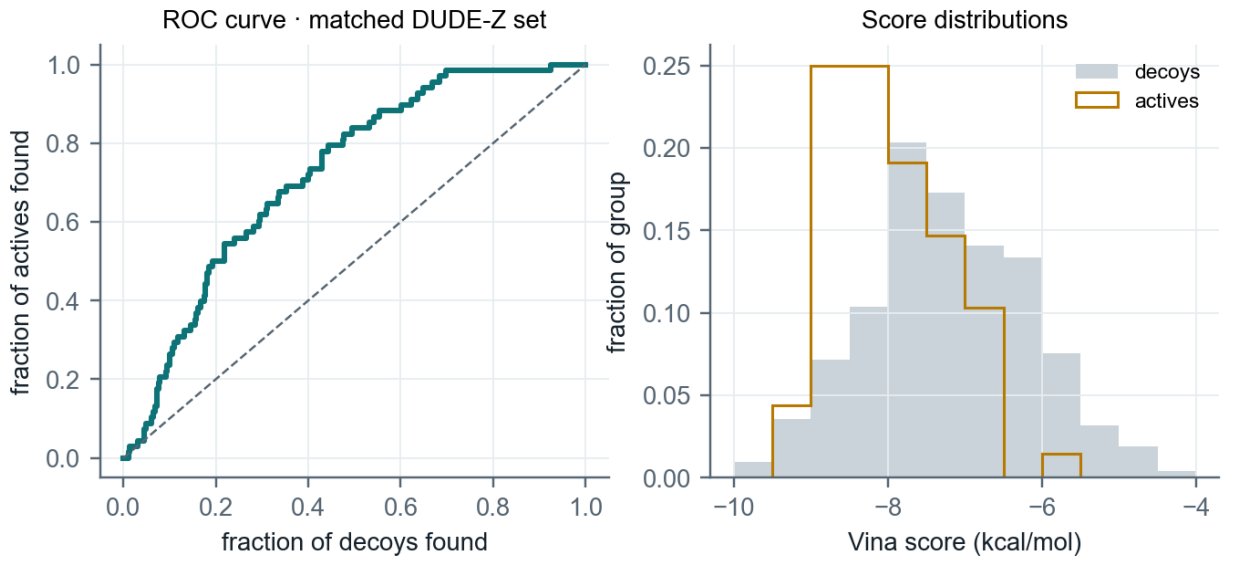


Figure 10.2 · The matched benchmark. 14 of 68 actives fall in the top 60 compounds (10%), about twice the random rate; none reach the top 1%. The active and decoy score distributions overlap heavily.

This is a realistic result for plain Vina on a kinase: clearly better than random (AUC 0.72, $EF_{10\%}$ 2.1) but with weak early enrichment. Buying the top 1% would still have found no actives. In a real project you would now try rescoring (Vinardo, GNINA), ensemble docking with several EGFR structures, or pharmacophore filters for the hinge H-bond, and measure whether each change improves $EF_{1\%}$ on this same benchmark before screening an unknown library.

KEY POINTS

Always redock the reference ligand into the exact receptor you will screen against, with several seeds.

Decoys are only fair against the actives they were matched to. Check that actives and decoys have similar size and charge.

A respectable AUC can hide useless early enrichment. Report $EF_{1\%}$ and $EF_{5\%}$ too.

Structural waters can decide whether docking works (Lab 6, Lab 7).

10.6 Pitfalls

- **Covalent inhibitors.** Afatinib, osimertinib, dacomitinib, neratinib and canertinib bind EGFR covalently through Cys797. Non-covalent docking scores only their reversible contacts, so their real potency is under-estimated.
- **Receptor conformation.** Lapatinib binds a different, inactive EGFR conformation (PDB 1XKK) from the active-like 1M17. Docking it into 1M17 asks it to fit the wrong pocket shape.
- **Size bias.** Larger molecules score better. Check ligand efficiency or compare within size bands.
- **Score over-interpretation.** Differences below ~1 kcal/mol are noise. Never rank by the second decimal.
- **Preparation errors.** One wrong protonation or tautomer rule silently affects the whole library.

10.7 Follow-up for docking hits

1. Visual inspection of the top-ranked poses (Lecture 8 checklist).
2. Rescoring or consensus scoring with a second function (Vinardo, AD4, GNINA).
3. Molecular dynamics of the complex and MM-GBSA/PBSA estimates; free-energy perturbation for close analogues.
4. ADMET prediction (e.g. SwissADME) to remove poor drug candidates.
5. Purchase and test in a biochemical assay. Only experiments establish activity.

Self-test

1. A screen of 2 000 compounds with 40 actives finds 6 actives in the top 20. Calculate $EF_{1\%}$.
2. What does a ROC AUC of 0.85 mean in plain words?
3. Give two reasons why an approved drug might rank poorly in a docking screen against its own target.

LECTURE 11

Publication figures and reporting

LEARNING OBJECTIVES

- Produce a clear 3D binding-site figure and a 2D interaction diagram.
- Write a reproducible docking methods section.
- Present docking results honestly.

11.1 3D figures

- Show the receptor as a light, semi-transparent cartoon; the binding-site residues (within ~4 Å) as thin sticks; the ligand as thicker sticks in a contrasting colour.
- Draw hydrogen bonds as dashed lines with heavy-atom distances (2.8–3.2 Å), not H···A distances.
- Label only the residues you discuss; move overlapping labels by hand.
- White background, ray-traced, at least 300 dpi at final print size (2400 px wide for a full-width figure).
- Use colours that survive greyscale and colour-blindness: avoid relying on red against green.

```
remove hydro
show cartoon, rec ; color grey80, rec ; set cartoon_transparency, 0.6
select site, byres rec within 4 of lig
show sticks, (site and not name N+C+O) or lig
util.cbaw site ; util.cbay lig
distance hb, lig, site, mode=2            # polar contacts, heavy atoms
label (byres rec within 3.5 of lig) and name CB, "%s%s" % (resn, resi)
orient lig ; zoom lig, 5 ; bg_color white
png figure_3d.png, width=2400, height=1800, dpi=300, ray=1
```

Lessons from testing the Lab 8 script: labelling every residue within 4 Å produced about 30 overlapping labels, and with hydrogens present PyMOL measured H···O distances (1.8–2.2 Å) that look wrong to readers. Removing hydrogens and labelling only residues within 3.5 Å fixed both (Figure 8.1).

11.2 2D interaction diagrams

2D diagrams show all interactions at once without perspective problems. Use Discovery Studio Visualizer (Receptor–Ligand Interactions → 2D Diagram), PLIP, LigPlot+ or PoseView. Build the complex with PyMOL, not `obabel -j`: in our test, Open Babel renumbered the protein so Met318 appeared as residue 94, which would put wrong residue numbers in a paper.

11.3 The methods section

A reader must be able to reproduce your docking from the methods alone. A template filled in with the real Lab 2 settings:

EXAMPLE METHODS PARAGRAPH

The crystal structure of the Abl kinase domain in complex with imatinib (PDB 1IEP, 2.10 Å) was obtained from the RCSB PDB. Chain A was retained; waters and heteroatoms were removed. Missing atoms and hydrogens were added at pH 7.4 with PDBFixer 1.x (OpenMM 8.6), and the receptor was converted to PDBQT with Open Babel 3.2.1. The ligand was protonated at pH 7.4 with Open Babel and prepared with Meeko 0.8.0 (7 rotatable bonds). Docking was performed with AutoDock Vina 1.2.7 in a box centred at (15.190, 53.903, 16.917) Å with dimensions 18.7 × 26.7 × 23.5 Å, exhaustiveness 8, 9 output modes, energy range 3 kcal/mol and random seed 42. The protocol was validated by redocking the crystallographic ligand (heavy-atom, symmetry-corrected RMSD 0.25 Å, computed with obrms). Interactions were analysed with PLIP 3.0.1 and figures rendered with PyMOL.

11.4 Results tables and honest wording

Compound	Vina score (kcal/mol)	LE	Key interactions	Note
Imatinib (reference)	−13.27	0.36	Met318, Thr315, Glu286, Asp381 (H-bonds)	Redocking RMSD 0.25 Å
Compound X	…	…	…	…

- Write 'predicted binding energy' or 'docking score', never 'binding affinity'.
- Say a compound is 'predicted to bind' or 'a candidate for testing', not 'an inhibitor'.
- Report the validation (redocking RMSD, enrichment) alongside the results.
- Cite the software: Vina (Trott & Olson 2010; Eberhardt et al. 2021), AutoDock4 (Morris et al. 2009), Open Babel, PLIP, PyMOL.

Self-test

1. List six settings a reviewer needs to reproduce a Vina docking.
2. Why should H-bond distances in figures be heavy-atom distances?
3. Rewrite this sentence honestly: 'Compound 7 is a potent EGFR inhibitor with a binding affinity of −11.2 kcal/mol.'

LECTURE 12

Capstone project guide

LEARNING OBJECTIVES

- Plan and carry out a complete docking study independently.
- Write it up as a short scientific report.

12.1 Choosing a target

Pick a target with a good holo structure and known actives. Suggestions (confirm the details on rcsb.org before starting):

Target	PDB	Ligand	Why it works well
Cyclooxygenase-2	5KIR	rofecoxib	Classic anti-inflammatory target; many actives in ChEMBL
Acetylcholinesterase	4EY7	donepezil	Alzheimer's drug target; deep, aromatic gorge; in DUDE-Z (ACES)
Estrogen receptor α	1ERE	estradiol	Well-defined hydrophobic pocket; many ligands
DPP-4	1X70	sitagliptin	Diabetes target; several drug classes to compare

12.2 Plan

1. **Validate**: redock the crystal ligand (Lab 2). Do not continue until RMSD ≤ 2 Å.
2. **Benchmark**: 10–30 ChEMBL actives plus property-matched decoys (DUDE-Z if your target is there). Report $EF_{1\%}$, $EF_{5\%}$ and AUC.
3. **Screen** (optional): a focused set of 100–300 compounds, e.g. natural products from PubChem or COCONUT.
4. **Analyse** the top 5–10: interactions, LE, visual inspection, comparison with the reference ligand.
5. **Report**: two pages plus figures.

12.3 Report structure and marking guide

Section	Content	Marks
Introduction	Target, disease relevance, aim of the study	10
Methods	Every setting needed to reproduce the work (Lecture 11 template)	25
Validation	Redocking RMSD; enrichment metrics with ROC curve	20
Results	Results table; one 3D and one 2D figure; interaction analysis	25
Discussion	What the results do and do not show; limitations; next experiments	15
Presentation	Clear writing, correct units, honest wording, citations	5

BEFORE YOU SUBMIT

Could another student reproduce every number from your methods section alone?

Did you report validation before results?

Did you avoid claiming activity from docking scores?

APPENDIX A

Command reference

Environment

```
conda create -n dock -c conda-forge python=3.11 vina meeko rdkit openbabel \
      pymol-open-source pdbfixer plip propka numpy -y
conda activate dock
```

Receptor

```
curl -O https://files.rcsb.org/download/1IEP.pdb
pymol -cq prep.pml                    # keep chain, remove waters and heteroatoms, save receptor.pdb
pdbfixer receptor.pdb --add-atoms=all --keep-heterogens=none --ph=7.4 --output=receptor_fixed.pdb
obabel receptor_fixed.pdb -xr -O receptor.pdbqt
```

Ligand

```
curl -o lig.sdf "https://models.rcsb.org/v1/1iep/ligand?auth_asym_id=A&label_comp_id=STI&encoding=sdf"
curl -o drug.sdf "https://pubchem.ncbi.nlm.nih.gov/rest/pug/compound/name/gefitinib/SDF?record_type=3d"
obabel lig.sdf -O lig_H.sdf -p 7.4
obabel lib.smi -O lib.sdf --gen3d
obabel lib.sdf -O lib_min.sdf --minimize --ff MMFF94 --steps 500
mk_prepare_ligand.py -i lig_H.sdf -o lig.pdbqt
mk_prepare_ligand.py -i lib_min.sdf --multimol_outdir ligs
```

Docking and analysis

```
vina --config box.txt --ligand lig.pdbqt --exhaustiveness 8 --seed 42 --out out.pdbqt
grep "REMARK VINA RESULT" out.pdbqt
mk_export.py out.pdbqt -s out.sdf
obrms -f lig_H.sdf out.sdf                # RMSD of every pose vs the crystal ligand
obabel out.sdf -O best.sdf -l 1           # first pose only
plip -f complex.pdb -t -p -o plip_out
propka3 receptor.pdb
```

AutoDock4

```
autogrid4 -p receptor.gpf -l receptor.glg
autodock4 -p ligand.dpf -l ligand.dlg
```

APPENDIX B

Glossary

Term	Meaning
Active	A compound with measured activity against the target.
altLoc	Alternate location: an atom modelled in two or more positions in a crystal structure.
Apo / holo	Structure without / with a bound ligand.
B-factor	Atomic displacement parameter; high values mean mobile or poorly defined atoms.
Blind docking	Docking with a box covering the whole protein.
Cross-docking	Docking a ligand into a receptor structure solved with a different ligand.
Decoy	A presumed inactive molecule used to benchmark a screen, usually property-matched to actives.
DFG motif	Asp-Phe-Gly loop of kinases; 'DFG-in' (active) and 'DFG-out' (inactive) states bind different inhibitors.
Enrichment factor (EF)	Hit rate in the top fraction of a ranked list divided by the random hit rate.
Exhaustiveness	Vina's number of independent Monte Carlo search chains.
Gatekeeper	Kinase residue controlling access to a back pocket (Thr315 in Abl).
Grid box	The region in which poses are searched.
Hinge	Kinase segment linking the two lobes; ATP-site inhibitors H-bond to its backbone.
HETATM	PDB record for non-standard residues: ligands, waters, ions.
Lamarckian GA	AutoDock4's genetic algorithm in which local-search improvements are written back to the genes.
Ligand efficiency (LE)	$-\Delta G$ divided by the number of heavy atoms.
PAINS	Pan-assay interference compounds: substructures that give false positives in assays.
PDBQT	PDB format with partial charges (Q) and AutoDock atom types (T); used by AutoDock and Vina.
pK_a	pH at which an ionisable group is half protonated.
Pose	A position, orientation and conformation of the ligand in the site.
Redocking	Docking a crystal ligand back into its own structure to validate a protocol.
Resolution	Level of detail in an experimental structure; lower numbers are better.
RMSD	Root-mean-square deviation between corresponding atoms of two poses.
ROC AUC	Area under the ROC curve: probability that an active ranks above a decoy.
Rotatable bond	Non-ring, non-terminal single bond whose torsion the search varies.
Scoring function	Function estimating how favourable a pose is.
SMILES	Line notation encoding a molecule's atoms, bonds, rings and stereochemistry.
Structural water	A conserved, buried water that mediates protein–ligand contacts.
Torsion tree	ROOT and BRANCH hierarchy in a ligand PDBQT defining its flexible bonds.
Virtual screening	Computational ranking of a compound library to select molecules for testing.

APPENDIX C

Answers to the self-tests

Lecture 1

1. Outputs: poses and scores. Components: a search algorithm and a scoring function.
2. Hit identification (virtual screening of large libraries) and hit-to-lead / lead optimisation (binding-mode hypotheses, rationalising SAR). It is cheap per compound and gives structural insight.
3. Any two: scoring error; receptor flexibility not modelled; wrong protonation or tautomer; poor solubility, permeability or stability; assay issues; the molecule binds but does not inhibit.

Lecture 2

1. Avoid it if possible: look for a higher-resolution holo structure of the same protein, or check the electron density; if none exists, use it cautiously and validate by redocking.
2. Residue name: columns 18–20; chain: 22; residue number: 23–26.
3. HIV protease is a homodimer. The single active site lies at the interface, built from residues of both chains (e.g. Asp25 and Asp25′).

Lecture 3

1. Remove: other copies of the protein, bulk waters, crystallisation additives (and the ligand, saved separately). Keep: structural waters, metal ions or cofactors in the site.
2. HIE: the hydrogen should be on Nε2 so that it donates an H-bond to the carbonyl oxygen.
3. $1 / (1 + 10^{7.4 - 7.9}) = 1 / (1 + 0.316) = 0.76$, so about 76% protonated (neutral).
4. Its empirical terms (steric, hydrophobic, hydrogen bond) depend only on atom types and distances; electrostatics are not an explicit term.

Lecture 4

1. SMILES → 3D (Open Babel --gen3d or RDKit) → protonate at pH 7.4 (Open Babel -p, Dimorphite-DL) → minimise (MMFF94) → torsion tree and PDBQT (Meeko).
2. Only the phenyl–piperazine C–N bond. The ring bond is in a ring, C–CH_3 is terminal, and the amide bond is treated as non-rotatable.
3. Docking keeps bond lengths and angles fixed; any strain or distortion from 3D generation would remain in every pose and distort the score.

Lecture 5

1. 24 / 0.375 = 64 points per side (already even): npts 64 64 64.
2. The search space grows with volume, so more independent search chains are needed to explore it thoroughly enough to find the global minimum.
3. Any two: literature or mutagenesis data; superimposing a homologous holo structure; pocket-detection software (fpocket, P2Rank, CASTp).

Lecture 6

1. 6 + 9 = 15 (3 translation, 3 rotation, 9 torsions).
2. Improvements found by local search are written back into the individual's genes and inherited, like Lamarck's inheritance of acquired traits.
3. Under-sampling: the search did not converge. Increase exhaustiveness and repeat with several seeds until the best score and pose are reproducible.
4. To approximate the conformational entropy lost when rotatable bonds are frozen on binding: more flexible ligands are penalised.

Lecture 7

1. $\sqrt{((1 + 1 + 1 + 9)/4)} = \sqrt{3} = 1.73$ Å.
2. Superimposition removes errors in position and orientation, which are exactly what docking must get right; the RMSD would look far too good.
3. A scoring failure: the correct pose was found but ranked third. Try rescoring (Vinardo, AD4, GNINA), check protonation and waters, or use consensus.

Lecture 8

1. Any four, e.g. H-bond 2.5–3.5 Å donor–acceptor; salt bridge ≤ 5.5 Å; hydrophobic C–C ≤ 4.0 Å; π-stacking centroid ≤ 5.5 Å; cation–π ≤ 6.0 Å; metal ≤ 3.0 Å.
2. The T315I gatekeeper mutation removes this H-bond (and adds bulk), causing clinical resistance to imatinib in chronic myeloid leukaemia.
3. A: 10.2/41 = 0.25; B: 8.1/25 = 0.32. B is the more efficient starting point, despite the weaker score.
4. $K_d = \exp(-10.0/0.593) = \exp(-16.86) \approx 4.7 \times 10^{-8}$ M ≈ 47 nM.

Lecture 9

1. e.g. 10^6 library → 5×10^5 after filters → docking → top 1% = 5 000 → inspection and clustering → 50 purchased → 2–5 confirmed hits.
2. Otherwise a method could separate actives from decoys using trivial properties such as size or charge, and the enrichment would not reflect real recognition of the binding site.
3. 20 000 × 15 s / 4 = 75 000 s ≈ 21 hours.

Lecture 10

1. Top 1% = 20 compounds. EF = (6/20) / (40/2000) = 0.30 / 0.02 = 15.
2. If you pick one active and one decoy at random, there is an 85% chance the active is ranked higher.
3. Any two: it is a covalent binder; it binds a different receptor conformation than the one used; preparation errors (protonation, tautomer); scoring-function bias; important waters or metals missing.

Lecture 11

1. Software and version; PDB ID and chain; preparation steps (pH, waters kept, protonation tool); box centre and size; exhaustiveness (or GA runs/evaluations); random seed; number of modes and energy range; validation RMSD.
2. Hydrogen positions are modelled, not measured; heavy-atom distances (≈ 2.8–3.2 Å) are what crystallographers and readers expect and can compare with experiment.
3. 'Compound 7 had a predicted binding energy (Vina score) of –11.2 kcal/mol and reproduced the hinge hydrogen bond to Met793; it is a candidate for experimental testing as an EGFR inhibitor.'

APPENDIX D

Software and references

Software	Version	Reference
AutoDock Vina	1.2.7	Trott & Olson, J Comput Chem 31:455 (2010); Eberhardt et al., J Chem Inf Model 61:3891 (2021)
AutoDock4 / AutoDockTools	4.2.6 / MGLTools 1.5.7	Morris et al., J Comput Chem 30:2785 (2009)
Meeko	0.8.0	github.com/forlilab/Meeko
Open Babel	3.2.1	O'Boyle et al., J Cheminform 3:33 (2011)
RDKit	2026.03	rdkit.org
PDBFixer / OpenMM	OpenMM 8.6	Eastman et al., PLoS Comput Biol 13:e1005659 (2017)
PROPKA	3	Olsson et al., J Chem Theory Comput 7:525 (2011)
PLIP	3.0.1	Schake et al., Nucleic Acids Res (2025)
PyMOL (open source)	3.x	Schrödinger, LLC
DUDE-Z	–	Stein et al., J Chem Inf Model 61:699 (2021)

Further reading

- Forli et al. (2016) Computational protein–ligand docking and virtual drug screening with the AutoDock suite. *Nature Protocols* 11:905.
- Lipinski et al. (1997) Experimental and computational approaches to estimate solubility and permeability. *Adv Drug Deliv Rev* 23:3.
- Hopkins et al. (2004) Ligand efficiency: a useful metric for lead selection. *Drug Discov Today* 9:430.
- Baell & Holloway (2010) New substructure filters for removal of pan assay interference compounds (PAINS). *J Med Chem* 53:2719.
- Mysinger et al. (2012) Directory of useful decoys, enhanced (DUD-E). *J Med Chem* 55:6582.

Versions of AutoDock4 and MGLTools are those distributed by Scripps; they were not exercised in the automated tests, because Lab 3 uses the AutoDockTools graphical interface.