



From Molecule to Medicine

AI in Drug Discovery & Development

How computation helps us decide what to test next.

Maseera Siddiqui • Class 2



Learning objectives

- 1 Describe the major stages of early drug discovery.
- 2 Explain where AI and computational methods can support target and compound prioritisation.
- 3 Understand virtual screening, QSAR and molecular docking at a beginner level.
- 4 Explain why docking score is not the same as experimental binding affinity.
- 5 Describe how computational models can support ADMET and toxicity prediction.
- 6 Understand drug repurposing, generative molecular design and lead optimisation.
- 7 Recognise why computational predictions still require experimental validation.



The challenge: finding one useful molecule

*Imagine a library of
1,000,000 compounds.*

*Would you test every one
in the laboratory?*

**1,000,000
compounds**



**Computational
prioritisation**



**Smaller
shortlist**




**Experimental
testing**

*A recurring theme today: PRIORITISATION, not
magic.*



Before AI: what does drug discovery involve?



A hit is a compound showing activity worth investigating.

A lead is a more developed compound series with a stronger overall profile.

A development candidate must balance activity with selectivity, ADME, safety and developability.



Where can AI and computation help?

Target discovery

Find disease-related patterns in biological data

Compound screening

Prioritize molecules from large libraries

Activity prediction

Estimate properties or biological activity

Structure-based work

Study possible ligand–target interactions

ADMET

Predict selected developability or safety endpoints

Optimization

Explore trade-offs among multiple properties



Step 1: identifying a drug target

What is a target?

A biological molecule or system whose modulation may influence disease — often a protein, receptor or enzyme.

What data might help?

Genomics
Gene expression
Proteomics
Pathway data
Clinical data
Scientific literature

Where AI may help

Find patterns, rank candidate targets, integrate large datasets and generate hypotheses for biological investigation.

A predicted target still needs biological validation.



A Simple Example: COX-2 and Inflammation

Disease context

Inflammatory processes involve the production of prostaglandins, which contribute to pain and inflammation

Target


COX-2 is an enzyme involved in prostaglandin synthesis and is an established target of anti-inflammatory drugs.

Drug-discovery question

Can computational methods help us identify and prioritize new molecules that may inhibit COX-2?



Step 2: Understanding the target's structure



Experimentally determined structures can come from methods such as X-ray crystallography or cryo-EM.

AI-based structure-prediction systems can provide useful structural hypotheses when appropriate.

A predicted structure should be assessed for confidence and suitability before using it in downstream modelling.



Protein structure prediction: where AI enters

The basic problem

A protein's amino-acid sequence contains information related to its three-dimensional structure, but predicting that structure accurately is difficult.

AI contribution

Modern deep-learning systems can predict protein structures and confidence measures from sequence and related information.

Why it matters

Structural information can support binding-site analysis, hypothesis generation and structure-based drug-discovery workflows.

Prediction is useful information — not a substitute for every experimental structure.



Virtual screening

A very large compound library

A manageable shortlist

**Compound
library**



**Computational
screen**



Rank / filter



**Candidates for
further study**

*Virtual screening = computational prioritisation
of compounds.*



Two broad virtual-screening strategies

Ligand-based screening

Starts from information about known active molecules.

We ask: which other molecules have relevant similarities or learned features associated with activity?

Useful when ligand information is available.

Structure-based screening

Uses three-dimensional information about the biological target.

We ask: how might candidate molecules fit or interact with the target's binding site?

Docking is one common structure-based technique.



QSAR: connecting structure with activity

*Chemical structure
+ molecular information*

QSAR model

*Predicted biological
activity*

**Known compounds
+ activities**



**Learn structure–
activity relationships**



**Predict activity of
new compounds**

*QSAR = Quantitative Structure–Activity
Relationship.*



What does the computer 'see' when it sees a molecule?

Molecular weight

How heavy is the molecule?

LogP / lipophilicity

How strongly does it prefer lipid vs water?

H-bond donors/acceptors

Potential hydrogen-bonding features

Polar surface area

A descriptor related to molecular polarity

Rotatable bonds

A simple measure related to flexibility


Fingerprints

Encoded structural features/substructures

Descriptors translate aspects of chemistry into computable information.



Molecular docking: the basic idea



*Traditional molecular docking is computational modelling
— not automatically AI.*




What is a binding pose?


Same ligand



Pose 1



Pose 2



Pose 3

Different possible orientations

Docking asks: how might this ligand orient itself in the binding site?



What interactions might we examine?

Hydrogen bonds

Directional interactions involving suitable donor/acceptor groups

Hydrophobic contacts

Non-polar regions interacting in a hydrophobic environment

Electrostatic interactions

Attraction or repulsion involving charged/polar groups

π interactions

Interactions involving aromatic systems

A plausible interaction pattern is useful — but still a model of reality.



Docking score $\neq$ experimental binding affinity

Docking score

Generated by a scoring function.

Useful for comparing or ranking poses/compounds within the assumptions of a docking workflow.

Its units and interpretation depend on the software and scoring function.

Binding affinity

An experimentally measurable thermodynamic concept describing the strength of ligand–target binding.

Examples of experimental quantities include K_d or K_i in suitable assays.

Do not report a docking score as though it were measured affinity.



A good binder is not automatically a good drug

*Excellent predicted
binding*

≠

Successful medicine

It may be poorly soluble.

It may not reach the target tissue.

It may be metabolized too rapidly.

It may interact with unwanted targets.

It may be toxic.

It may be difficult to formulate or manufacture.



ADMET: what happens beyond target binding?

A**Absorption**

Can the drug reach systemic circulation appropriately?

D**Distribution**

Where does it go in the body?

M**Metabolism**

How is it chemically transformed?

E**Excretion**

How is it eliminated?


T**Toxicity**

Could it cause harmful effects?

A promising molecule must survive more than a docking experiment.



AI in ADMET and toxicity prediction



Possible endpoints include solubility, permeability, metabolic liabilities, transporter interactions or selected toxicity outcomes — depending on the model.

Different endpoints require different datasets and validation strategies.

An in silico prediction is a screening aid, not a declaration that a compound is safe.



Drug repurposing: a new job for an existing drug?

The idea

Drug repurposing seeks a new therapeutic use for an existing or previously investigated compound.

How AI may help

Analyse relationships among drugs, targets, diseases, genes, pathways, literature and clinical data to generate new hypotheses.


The catch

A computationally suggested new indication still requires biological evidence and, where relevant, clinical evaluation.

Repurposing can shorten some early steps — it does not remove the need for evidence.



Generative AI and de novo molecular design



Generative models can propose new molecular representations rather than only rank existing molecules.

Researchers can use property objectives or constraints to guide generation.

Generated molecules still need chemistry checks, synthesis feasibility assessment and experimental testing.



Lead optimisation is a balancing act

Potency

Does it produce the desired biological effect at useful concentrations?

Selectivity

Does it avoid unwanted targets?

ADME

Can it reach and persist at the right place for the right time?

Safety

Are toxicity liabilities acceptable for further development?


Developability

Can it be synthesised, formulated, stored and manufactured?

The 'best' molecule is rarely the one with only the highest potency.



Putting the pieces together



Illustrative only: real projects may change the order, repeat steps, or use different tools depending on the question.

Computation helps decide where to spend experimental time and resources.



What can go wrong?

Poor data

Wrong labels, missing data or narrow chemical space

Overfitting

Excellent training performance, poor generalisation

Wrong structure/context

Protein conformation or binding site may be unsuitable

False positives

A promising prediction may fail experimentally

Limited applicability

Model may be used outside the domain it learned

Overinterpretation

Treating scores as proof rather than hypotheses

The solution is not 'avoid computation' — it is validate and interpret it properly.



Class activity: build a screening pipeline

You have 100,000 compounds and want to find potential inhibitors of Protein X.

Your available tools

ML / QSAR • molecular docking • property filters •
ADMET prediction • experimental assays

Arrange the tools into a sensible screening strategy.

For each step, say what information goes in and what comes out.

Where would you reduce the number of compounds?

Which results would you NOT treat as proof?

Why shouldn't you simply choose the compound with the best docking score?

There is no single universal pipeline — justify your choices.



Key takeaways

1

AI can help researchers navigate enormous chemical and biological datasets.

2

AI is especially useful for prioritising targets and compounds — not magically producing a finished medicine.

3

QSAR, molecular docking, ADMET prediction and AI are related computational tools, but they are not interchangeable terms.

4

A strong docking score or AI prediction does not establish efficacy, affinity or safety.

5

Drug discovery is multi-objective: potency must be balanced with selectivity, ADME, safety and developability.

6

Computational predictions ultimately need appropriate experimental — and for medicines, clinical — validation.

From millions of possibilities → better-informed experiments.

