

# VANTAI

## ILUMINATE THE MOLECULAR WORLD

Neo-1 | 2025

# Computational approaches for small molecule drug discovery

## Introduction: computers and drugs

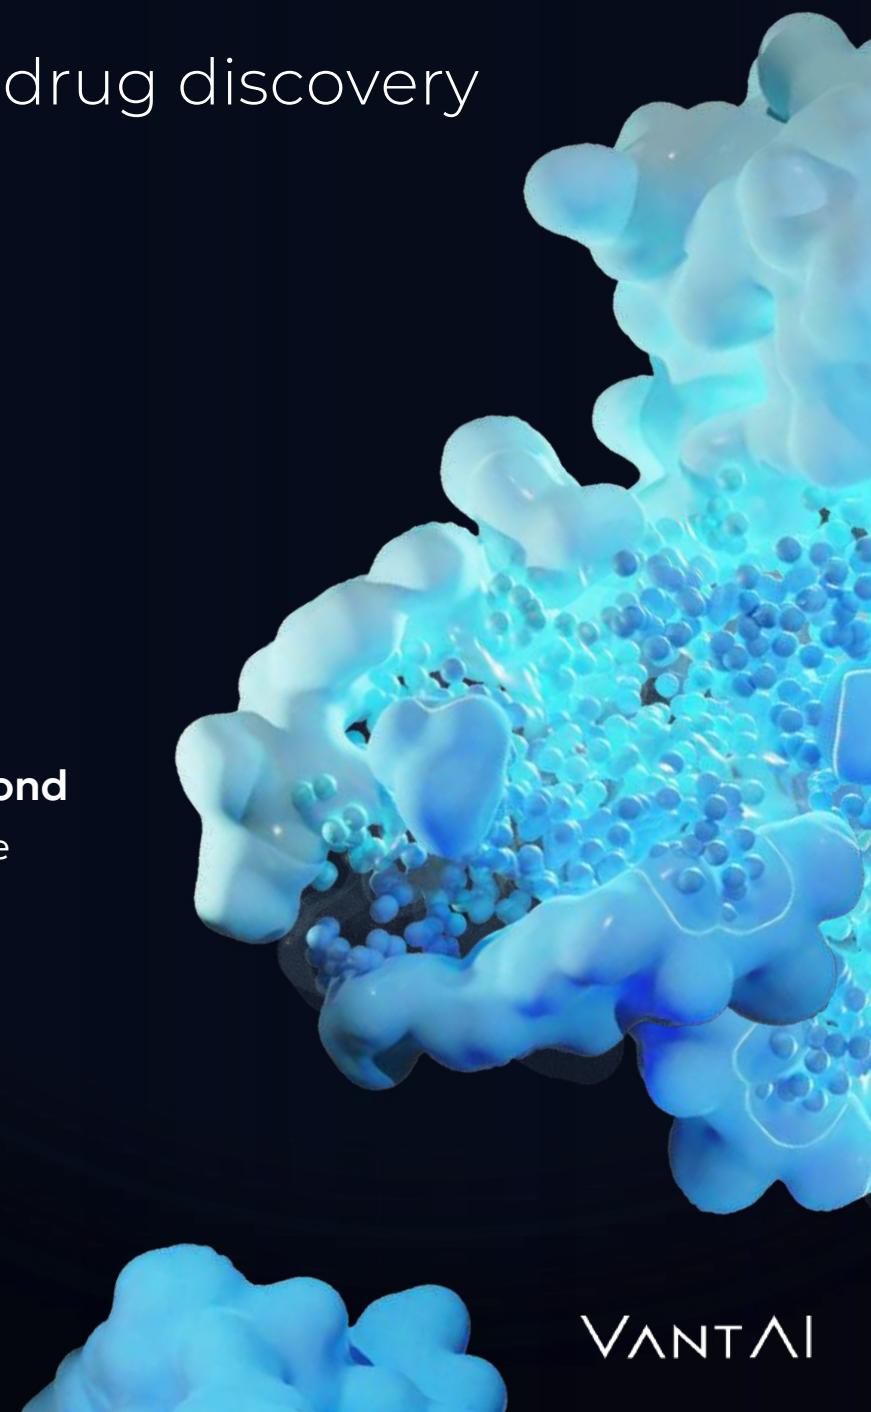
- Introduction
- The problem: the art and science of making drugs
- A quick history of computational drug discovery

## The “classic” era: task-specific tools and models

- Single task focused methods
- Breakthrough in protein structure prediction: the road to AF2 and beyond
- De-novo molecular generators

## The new era: foundation models trained on black box data – Neo and beyond

- Neo-1: unifying all-atom structure prediction and de-novo generation for the first time
- The promise of black-box data & NeoLink



# Who are we?



Luca Naef  
Co-Founder &  
CTO

- BSc/MSc ETH Zurich - Molecular Bio & Deep Learning
- Research in Stanford, UNSW, TokyoTech
- Software Engineer & first tech startup during BSc
- Diverse roles in Biotech - Regeneus (AUS), CJ Partners (JPN)
- QuantumBlack & McK - AI in Drug Discovery across Fortune 100/500
- Co-founded VantAI in 2019



Vladas Oleinikovas  
Director of Comp Chem

- BA/MSci Cambridge - Nat. Sci. / Chem.
- PhD UCL - Chemistry
- Senior Scientist at UCB Pharma
- Acting Head of CADD at Monte Rosa Tx
- Co-inventor of clinical VAV1 degrader
- Joined VantAI in 2024

# Finding new medicines is hard

~1-3B

Cost to find a new drug<sup>1</sup>

90+%

Chance of failure **after** entering human trials<sup>2</sup>

13+ yrs

Typical development **after** cause of disease identified

1. Estimates vary – e.g. from 0.88B (Eastern Research Group, “Drug Development Final Report”, Sept. 2024, for U.S. Department of Health and Human Services) to 2.6B (DiMasi et al., J Health Econ 2016)

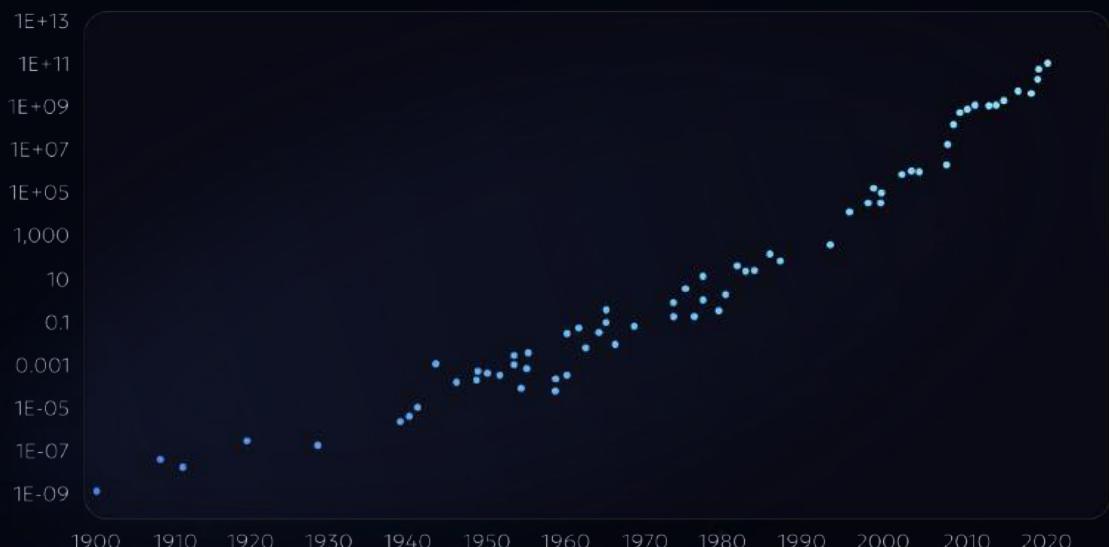
2. Smietana et al., Nature Reviews Drug Discovery, 2016

3. Paul et al., Nature Reviews Drug Discovery, 2010

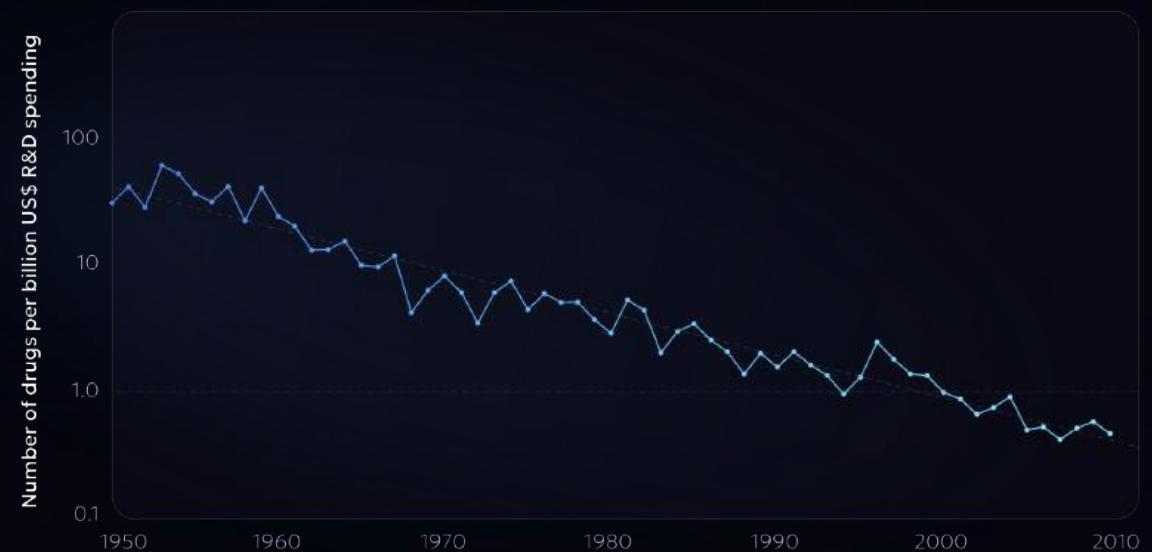


And has been getting harder

Moore's Law: Transistors

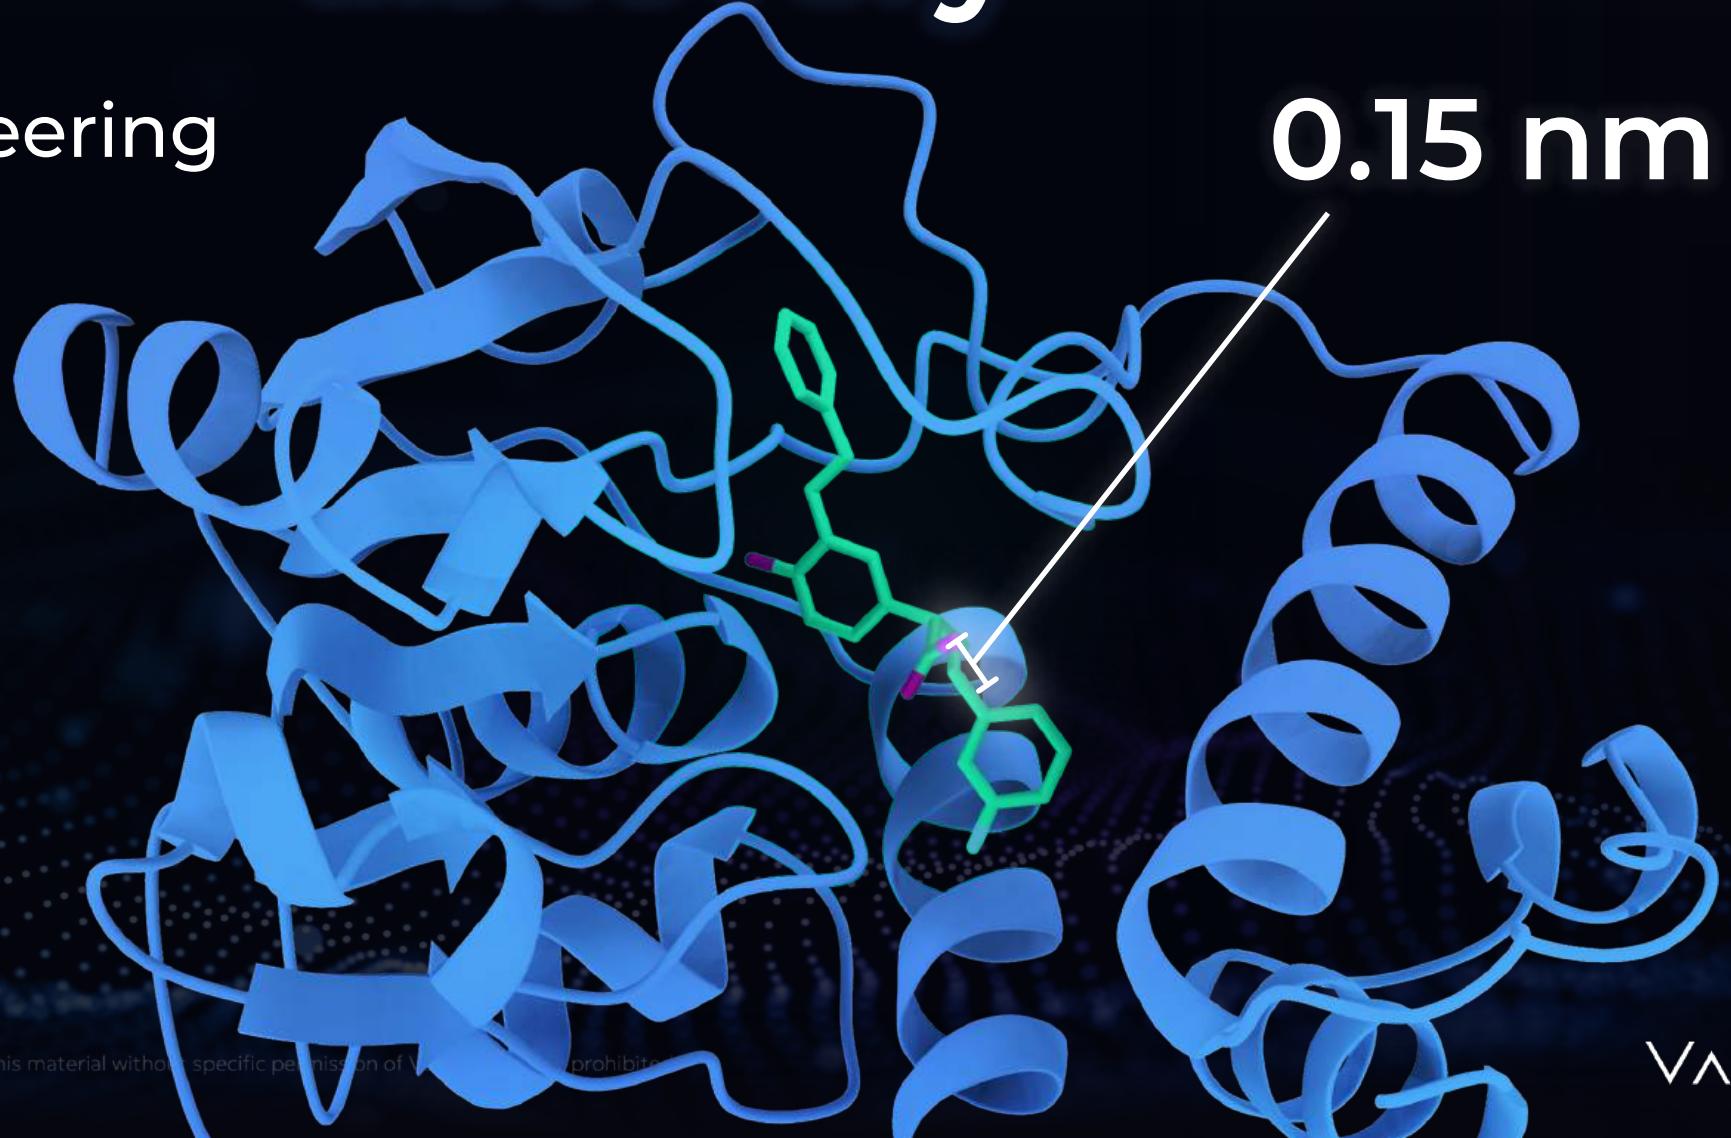


Eroom's Law: Medicines

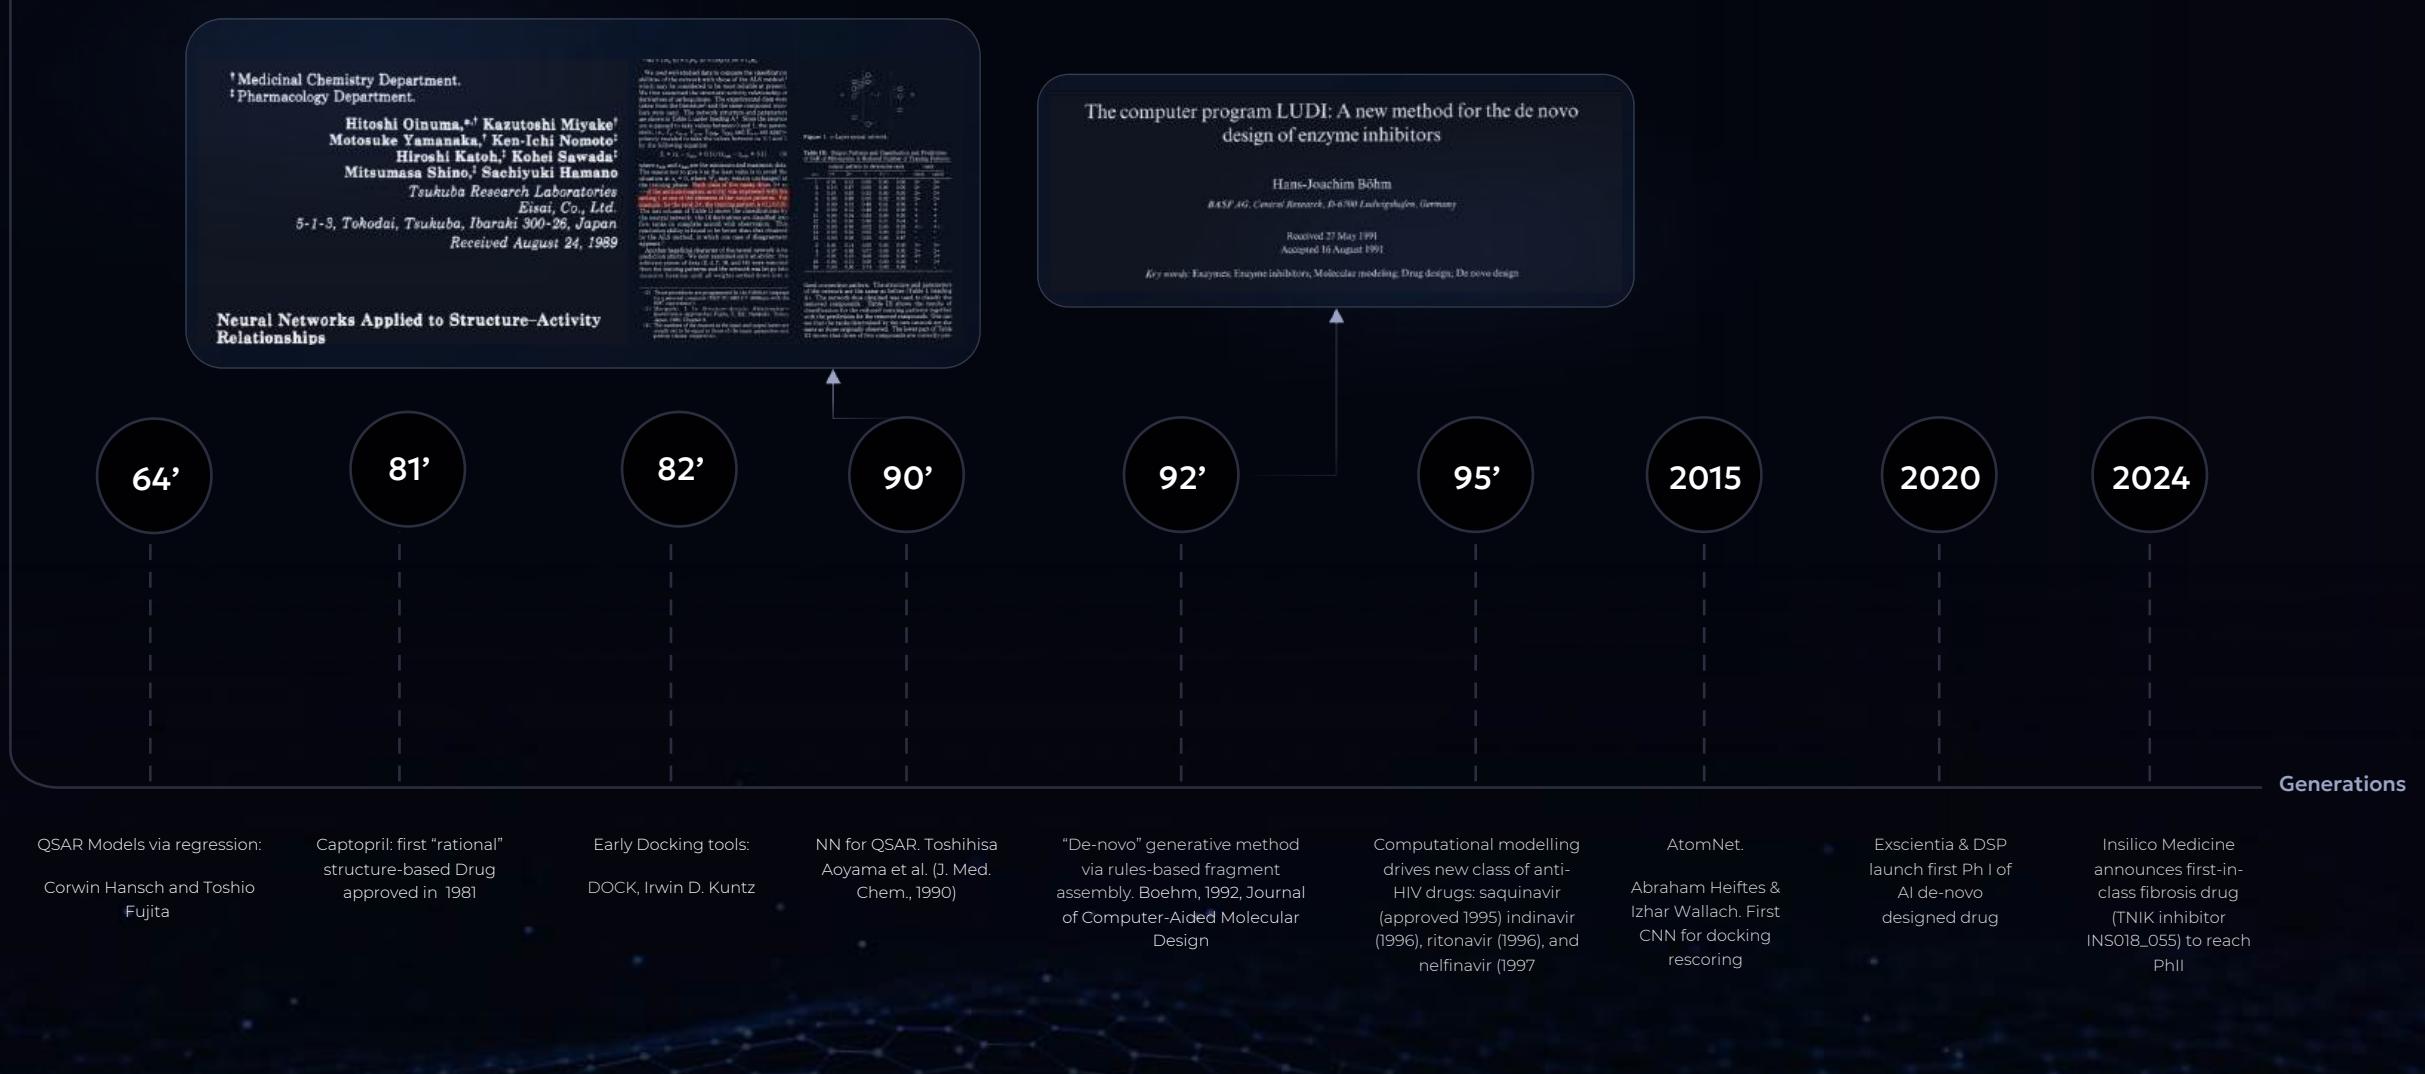


# Goal: Rational Drug discovery

Precision  
nano-engineering  
of therapies



# This is, by no means, a new idea!



# 1981



Xu, The path to the next computational transformation of drug discovery, Medium, 2022

However, we should not forget that **irrational** drug discovery is incredibly successful



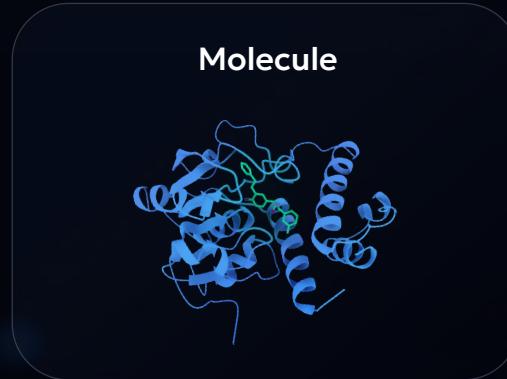
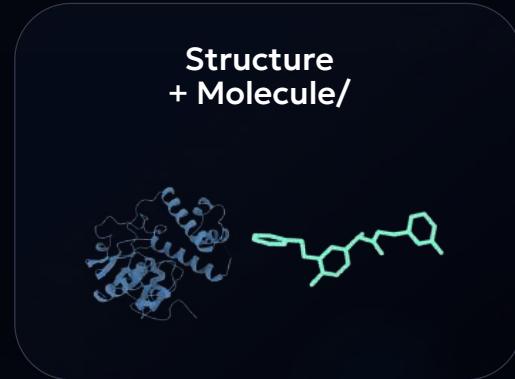
haloperidol  
1958



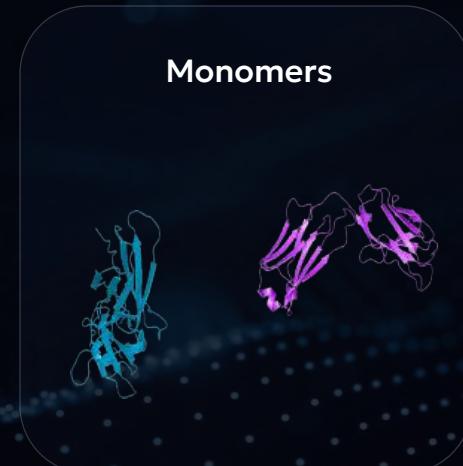
”It was a matter of life or death — a matter of survival. I **don't believe we really had a clear-cut strategy**. We were simply doing whatever we could, and there weren't many things that we could do then. We didn't have much money and there were not many researchers. **We had to make a lot of simple compounds as quickly and possible and screen them using very simple methods**”

Paul Janssen

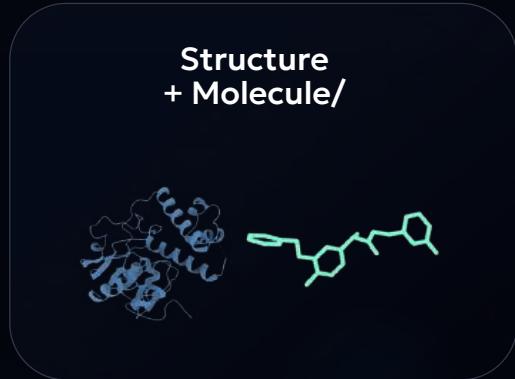
# The classical modelling tools focus to solve one task



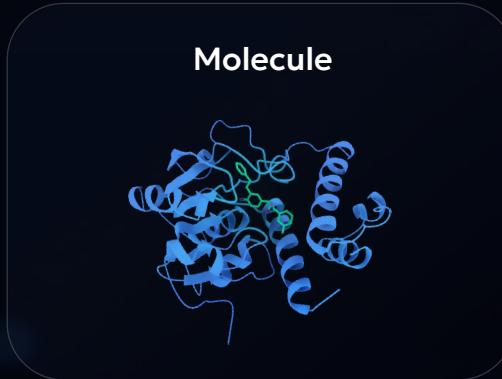
- Optimized to solve one task efficiently, eg. docking: sampling + scoring
- Relies on specific (often rigid) input, that may be prohibitively expensive (eg. X-ray structure) and limited by applicability



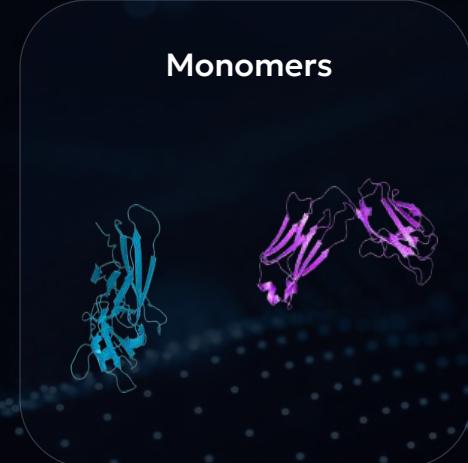
# The classical modelling tools focus to solve one task



Small molecule docking



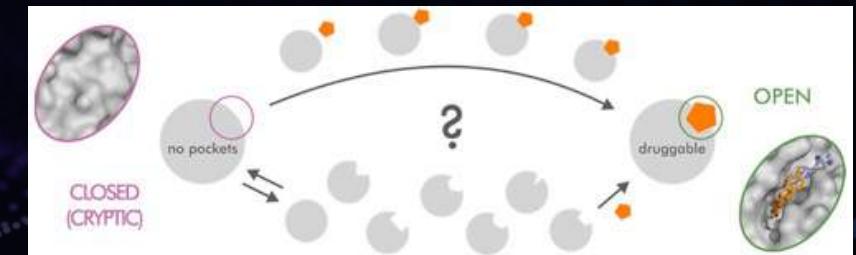
- Optimized to solve one task efficiently, eg. docking: sampling + scoring
- Relies on specific (often rigid) input, that may be prohibitively expensive (eg. X-ray structure) and limited by applicability



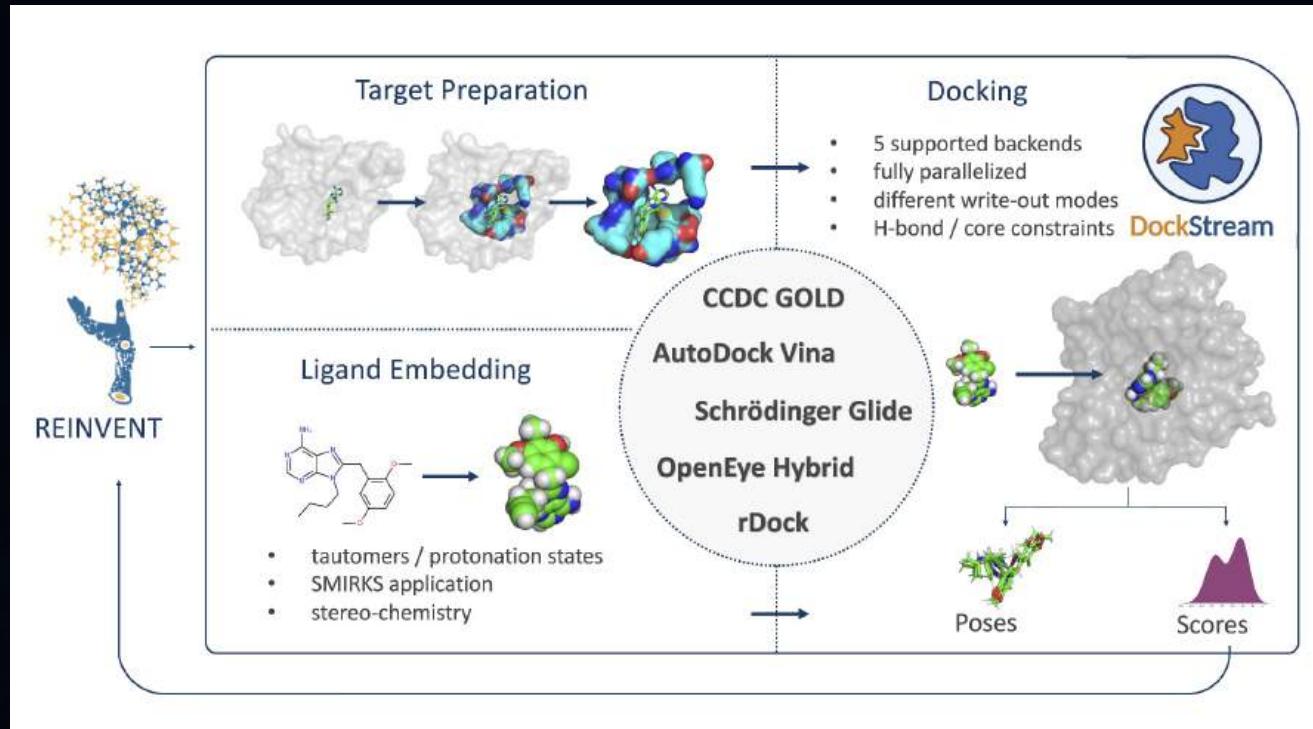
Protein-Protein  
Docking



- Scaling with Moore's law:
  - ✓ run more iterations, larger libraries / systems
  - ✗ sampling limited by inputs (esp. receptor flex.)



# Complex methods limited to sequential combination of tasks



<https://github.com/MolecularAI/DockStream>  
<https://pubs.acs.org/doi/10.1021/acs.jcim.0c01451>

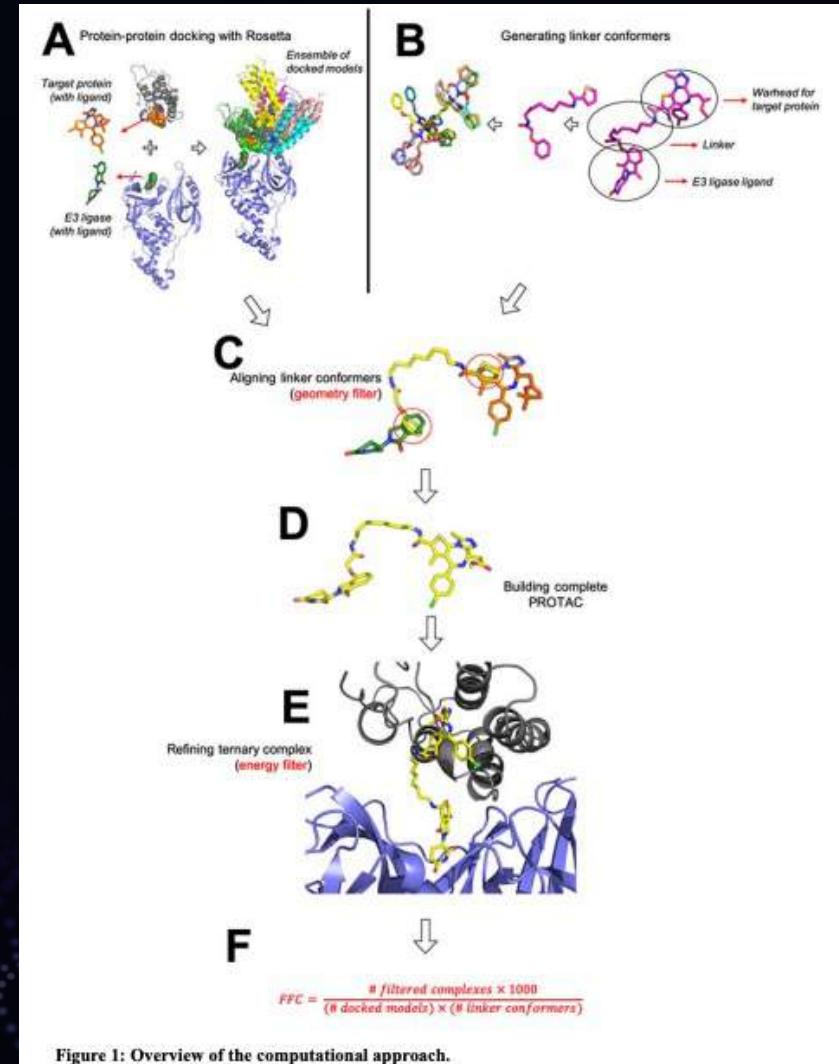
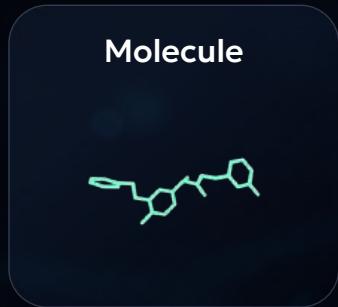


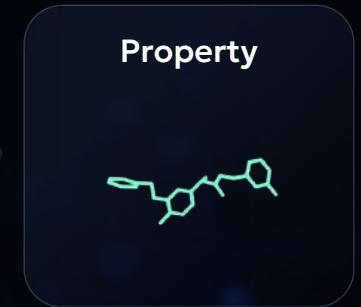
Figure 1: Overview of the computational approach.

There's broadly two flavors of machine learning that are used

## Discriminative

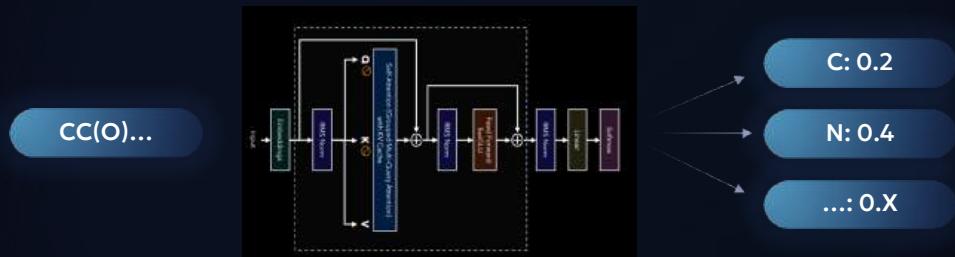


## Generative



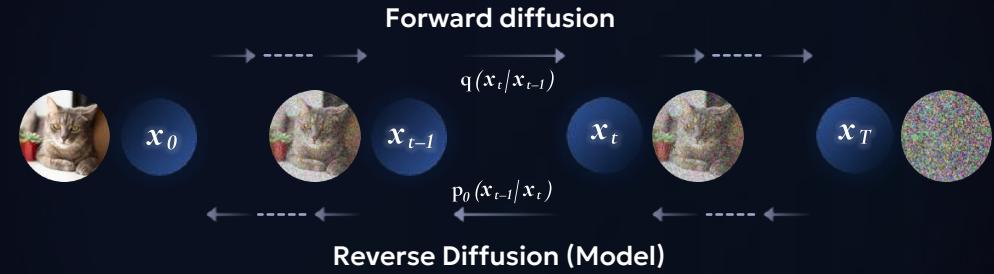
Two methods have been particularly critical  
in the emergence of generative methods

## Language models



- Trained to predict token: either next token (e.g. ChatGPT, “autoregressive”) or some masked tokens in middle of sequence (e.g ESM3)
- Breaks down problem into simpler steps: one token at a time
- Tokens require discrete data
- Became SOTA for language  $\approx$  with GPT-1 in 2018

## Diffusion models



- Trained to remove simple noise added during training process
- Breaks down problem into simpler steps: only predict “a bit” of noise – makes problem easier since e.g. high and low noise steps often very different problems
- Usually used for continuous data: images, coordinates, expression levels, ...
- Usually a lot more fancy math and many extensions (Flows, Schrodinger Bridges, ...)
- Became SOTA for images  $\approx$  with ADM in 2021

# Diffusion: A breakthrough approach for generative modelling

- Recent breakthroughs leverage Diffusion models
- Inspired by physical diffusion processes (brownian motion)
- Unsupervised training - stepwise noise addition to ground truth
- Model learns going from noised sample to ground truth
- Noise depends on problem at hand – Gaussian noise on pixels, latent noise or more complex (SO(3)-subspace diffusion)

Generative Adversarial Networks  
(2015)

Images of flowers



GANs & Roses

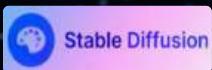
Diffusion Models  
(2020)



“Show a molecular glue holding together two planets”

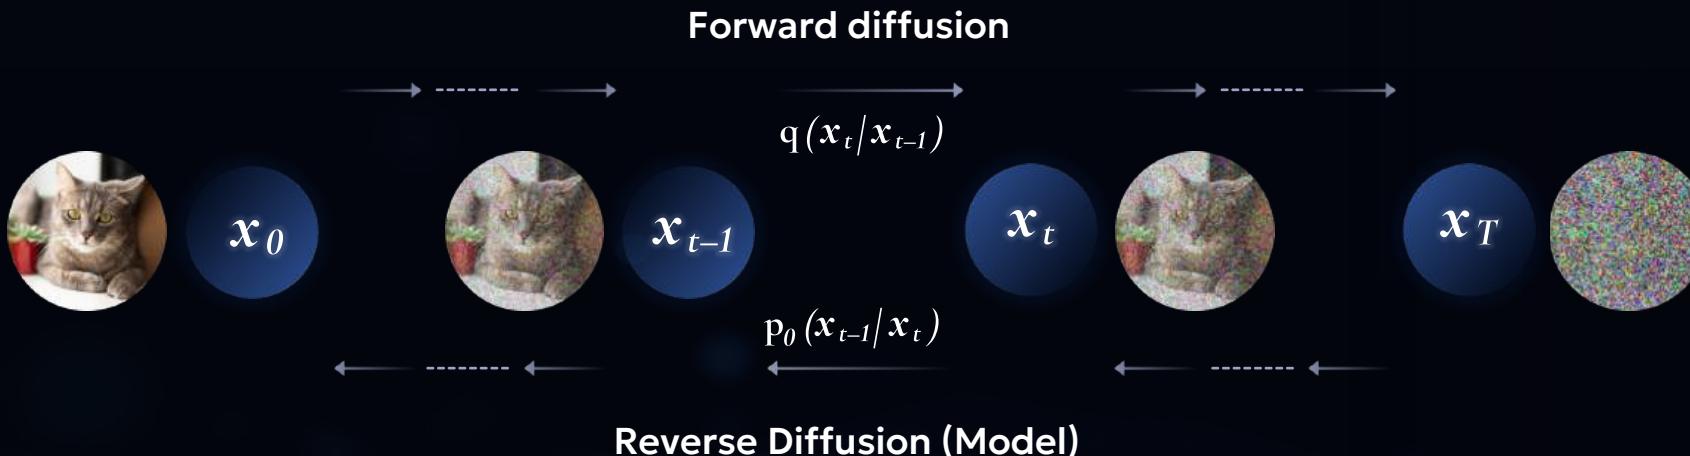


DALL-E 3



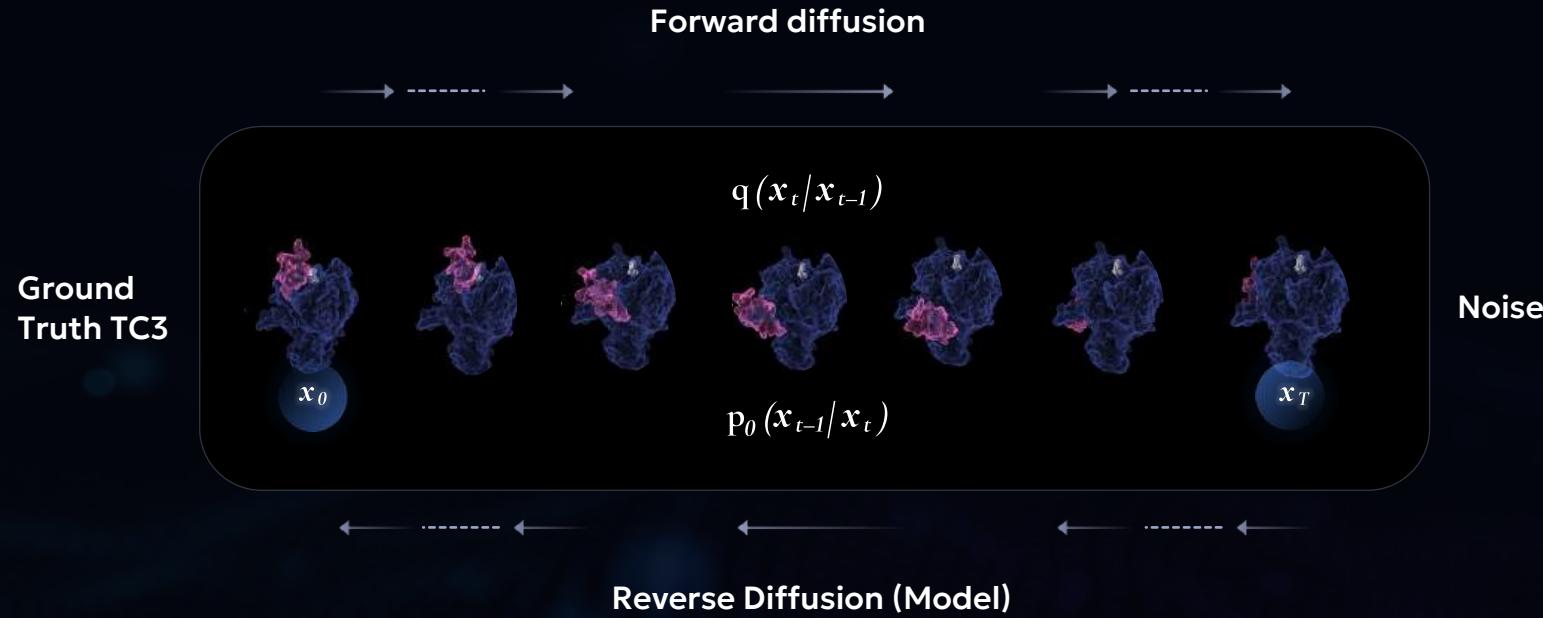
Imagen

Diffusion Models are trained unsupervised by adding noise to training data and predicting ground truth



Diffusion probabilistic models - Jascha Sohl-Dickstein, Google Brain Talks  
<https://www.youtube.com/watch?v=XCUIInHP1TNM>

In molecular modelling, “noise” can be added in highly flexible ways

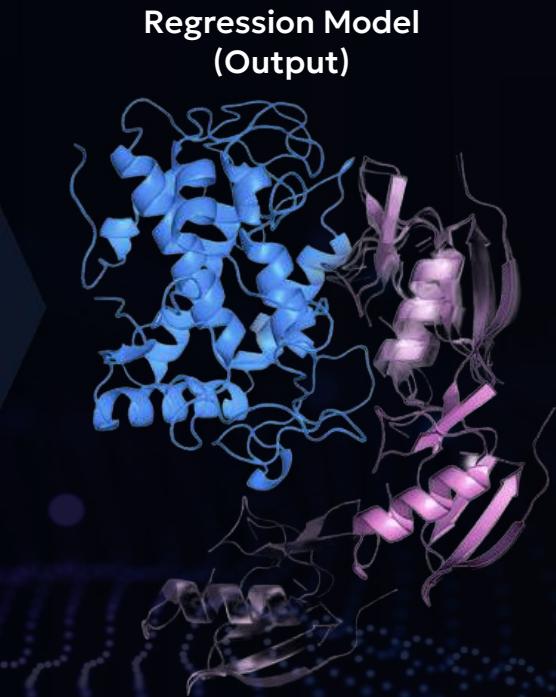
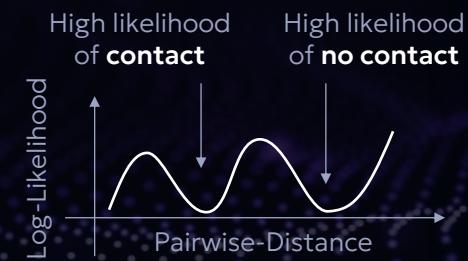
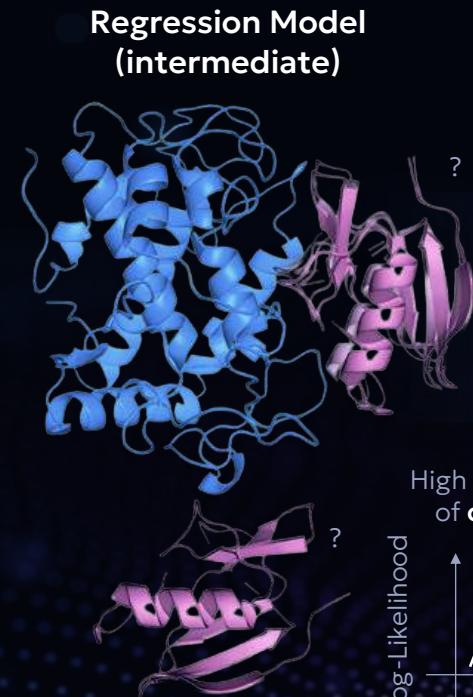
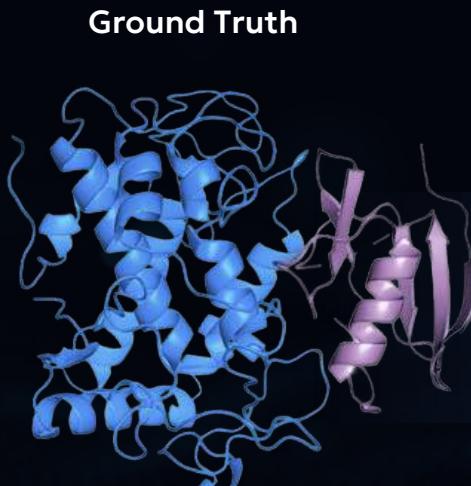


### Noise types

- Example: SE(3) diffusion (rotation, translation): **DiffMaSIF**
- Euclidean diffusion: **ApolloDiff**
- Latent Diffusion: **LatentDiff**

# Why use Diffusion for structure prediction problems?

Apart from the fact the protein structures are not static,  
Regression models (e.g. AF2) have a major issue

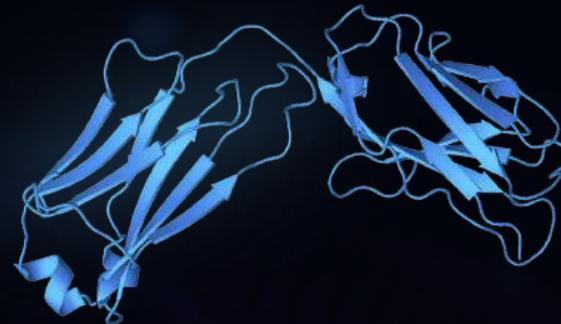


AI is generally used to predict one or multiple of the following given one or multiple of them as input

## Sequence

MGRLQLVVLGREDAHFIYENKDVSQ...

## Structure

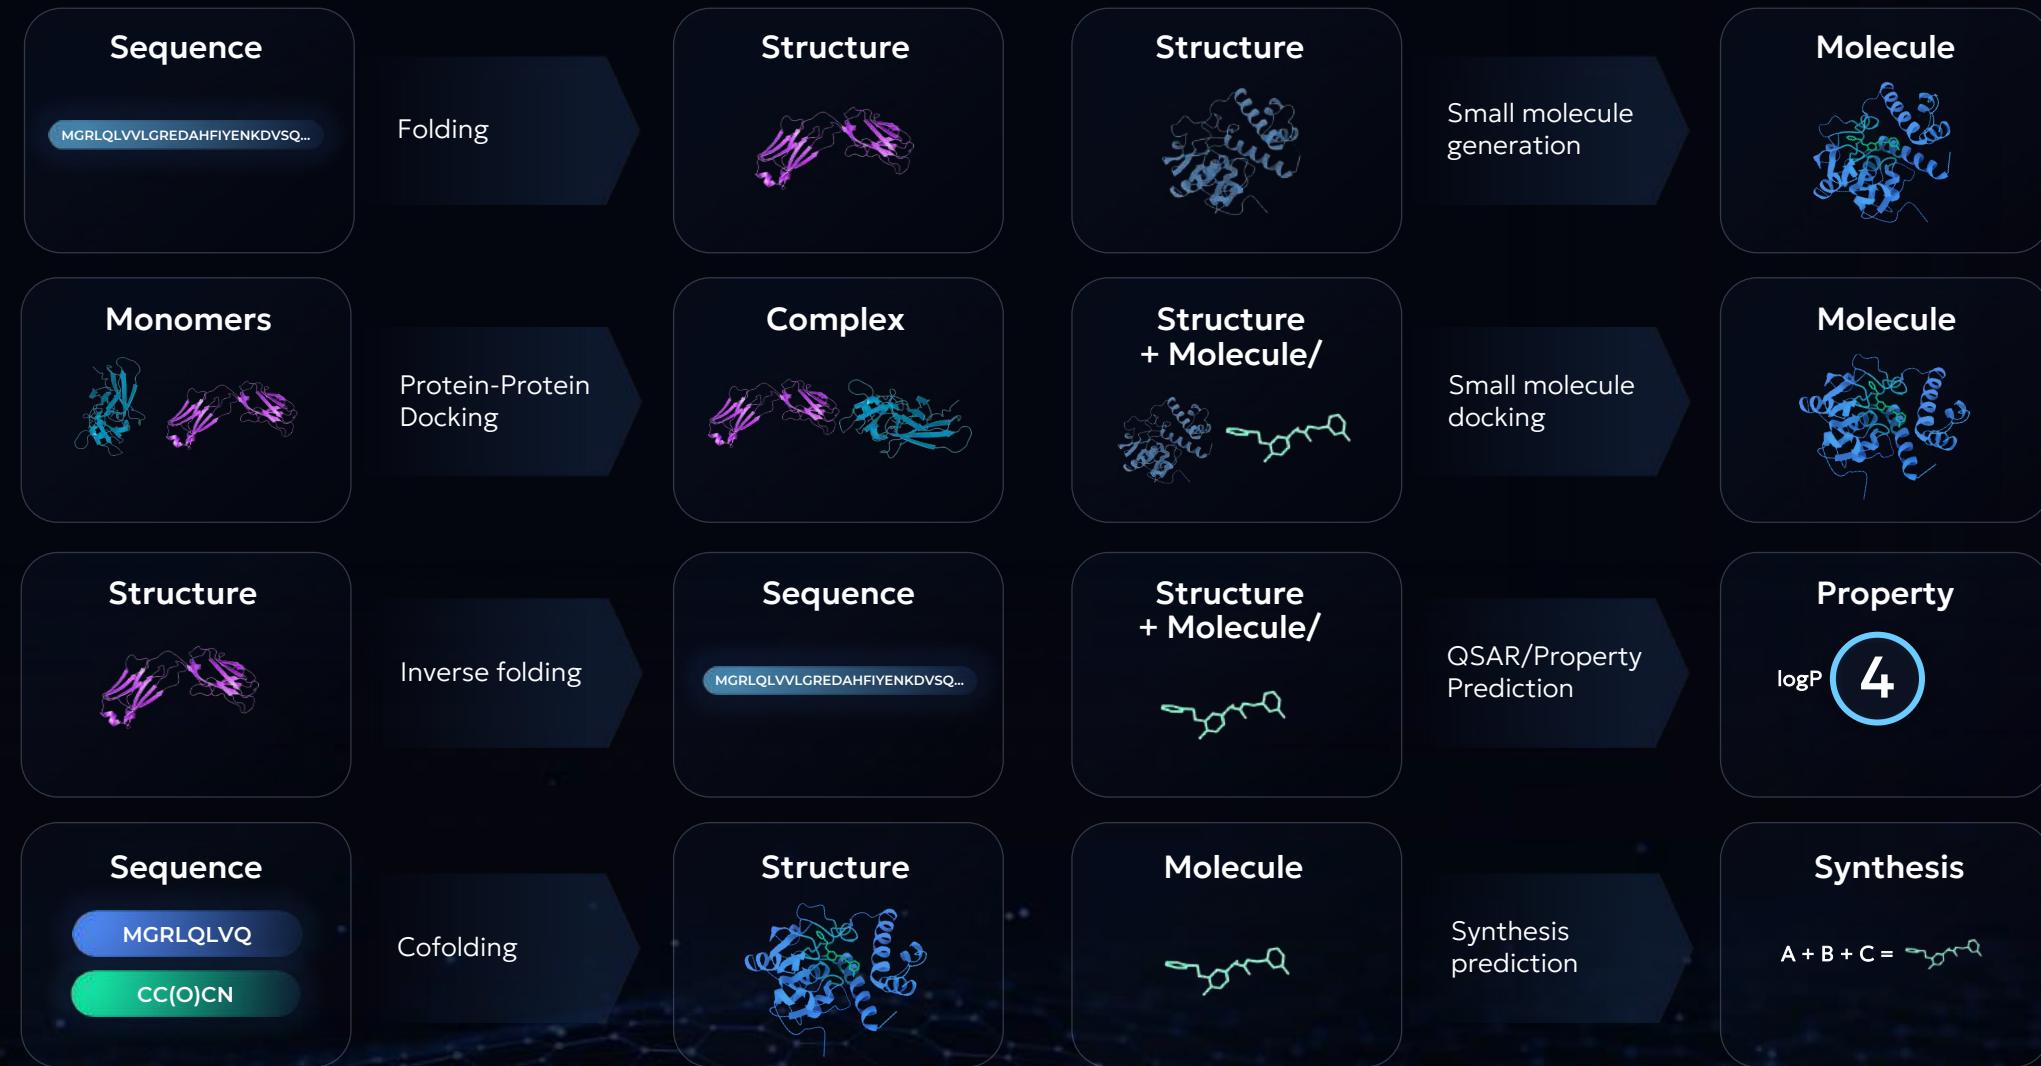


## Property

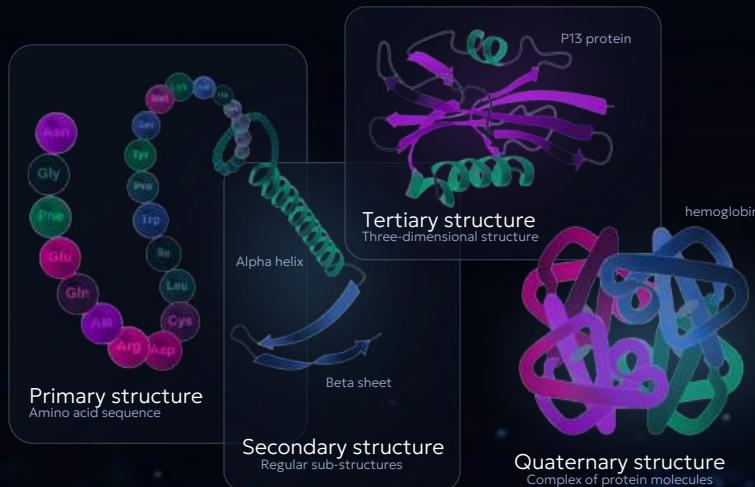
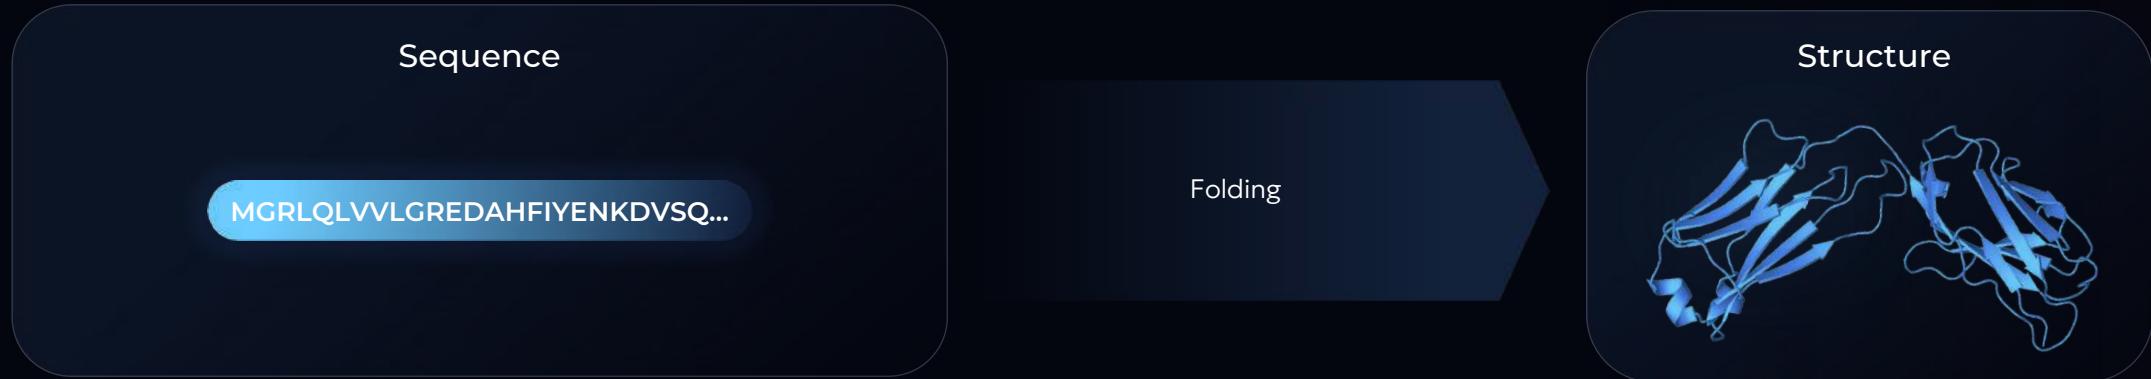
Kinase

Early ML methods mostly followed the classic methods focusing on the same tasks following the same limitations...

# This has led to a large "Zoo" of models in use



# Breakthrough: Protein folding



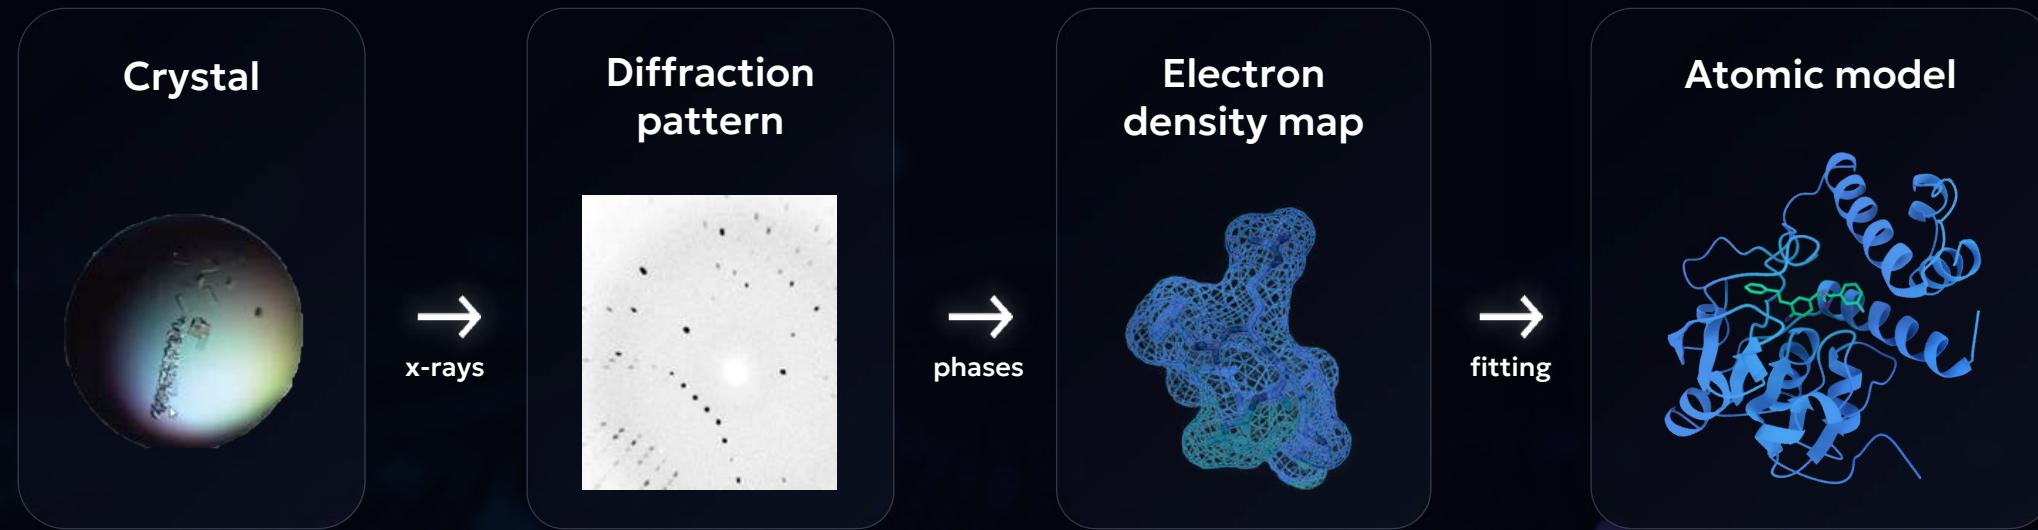
**A protein's structure is uniquely determined by its sequence**

**“Anfinsen’s Dogma”**  
– 1972 Nobel Prize



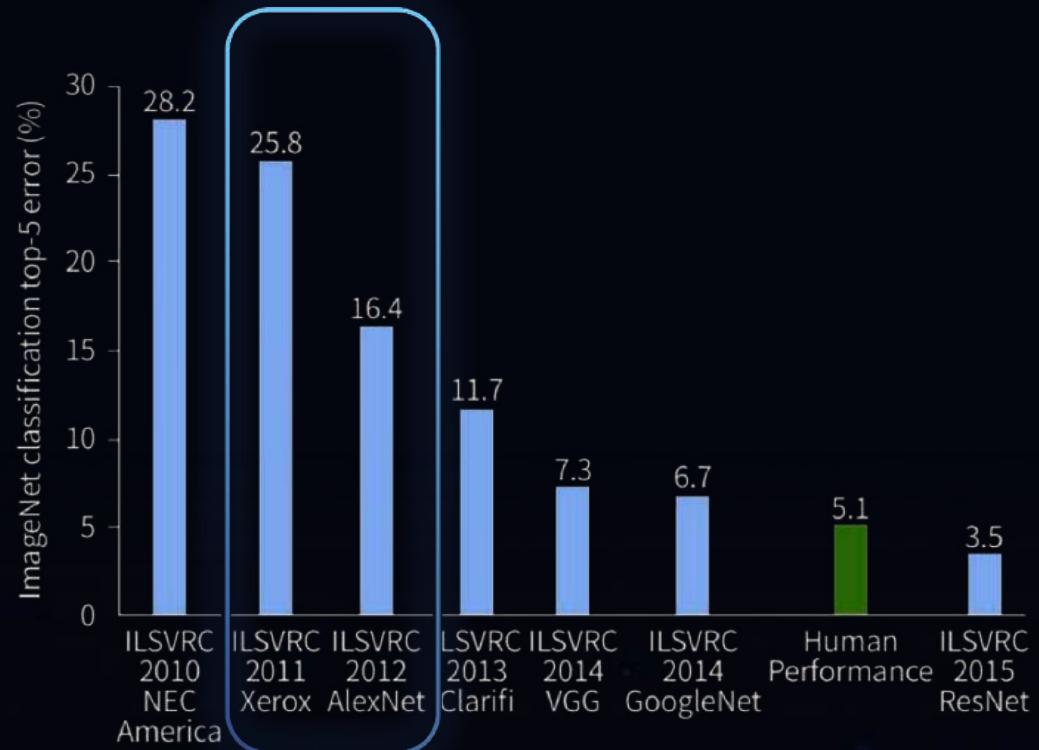
# Traditional route:

## X-ray crystallography



Up to **1M USD**  
Can take **years**

# AF2 was an “AlexNet” like moment



Median Free-Modelling Accuracy

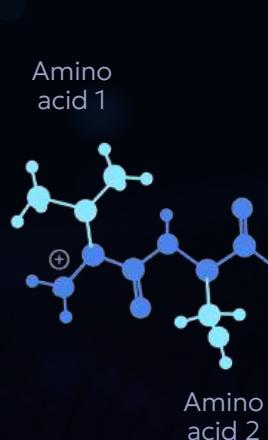


# Proteins are polymers comprised of amino acids

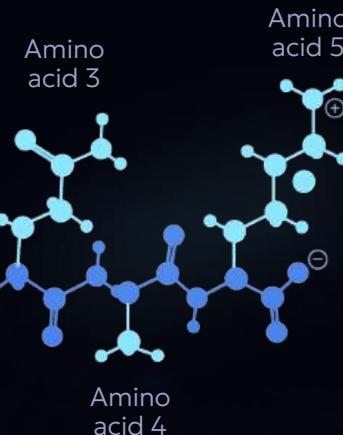


>UniRef90\_P02057 Hemoglobin subunit beta-1/2 n=6  
Tax=Euarchontoglires TaxID=314146 RepID=HBELRABIT  
MVHLSSEEKSATLWGKVNVVEVGGEALGRLLVYPWT  
QRFFESFGDLSSANAVMNWKVKAHGKKVLAFAEGLSHL  
DNLKGTFAKLSELHCDKLHVDPENERLLGNVLVIVLSSHFG  
KEFTPQVQAAQKVVAGVANALAHKYH

## Backbone

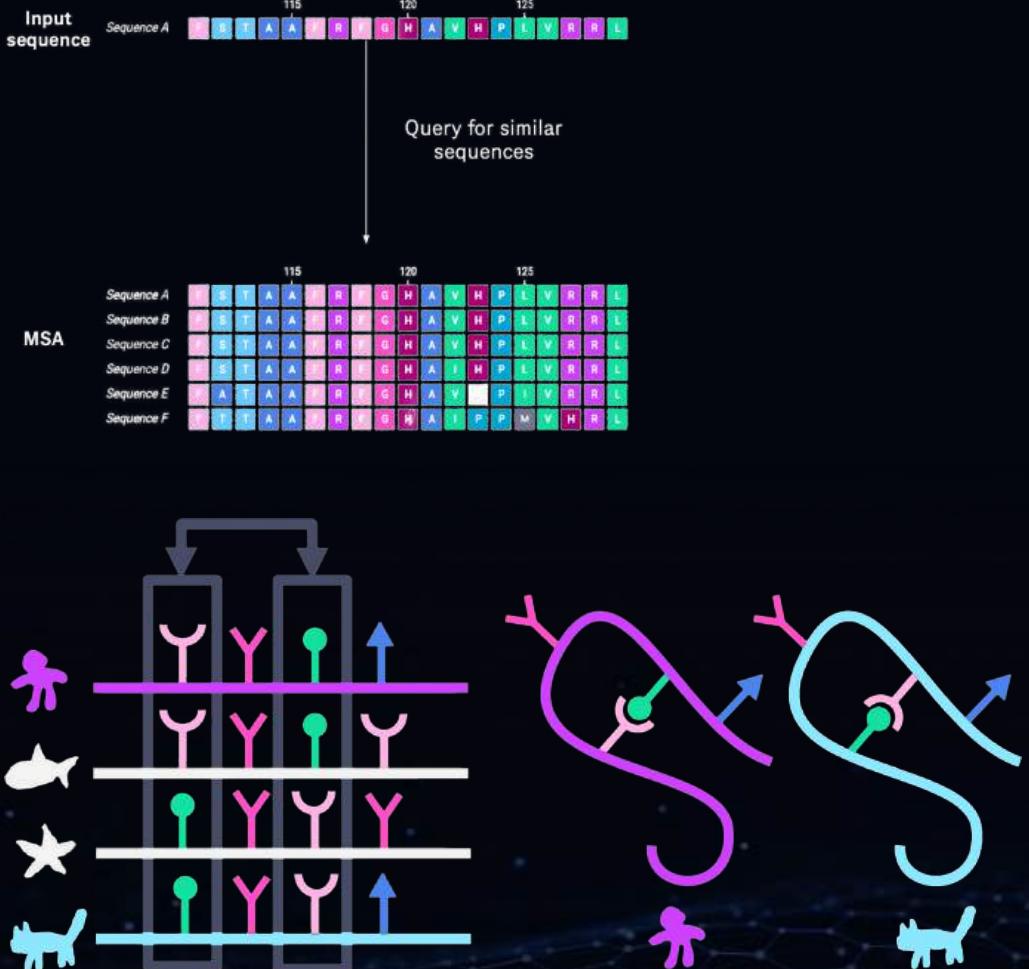


## Side chain



- We usually split structures in backbone (mainchain) & sidechain
- Backbone structure often assumed +/- independent of specific sidechains
- E.g. for protein design: often design backbones and then predict different residues that could take this shape

However, technically AlphaFold2 doesn't fold protein sequences, it predicts structures based on Multiple Sequence alignment



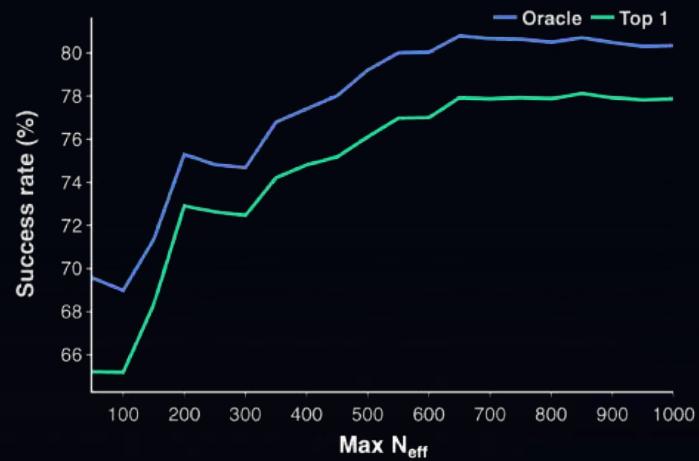
- MSA gives us **co-evolution**
- Residue pairs that co-evolve (e.g. charge switch on one always observed with opposing charge switch on other) are likely in contact
- Often ignored: it's also a form of **RAG** (**retrieval augmented generation**) – giving large number of sequences that likely fold into same structure makes sequence-based fold retrieval more robust

# In fact, co-evolutionary information is absolutely critical



Source: Zeming Lin  
et al., 2021

### AF2 PINDER performance by MSA depth

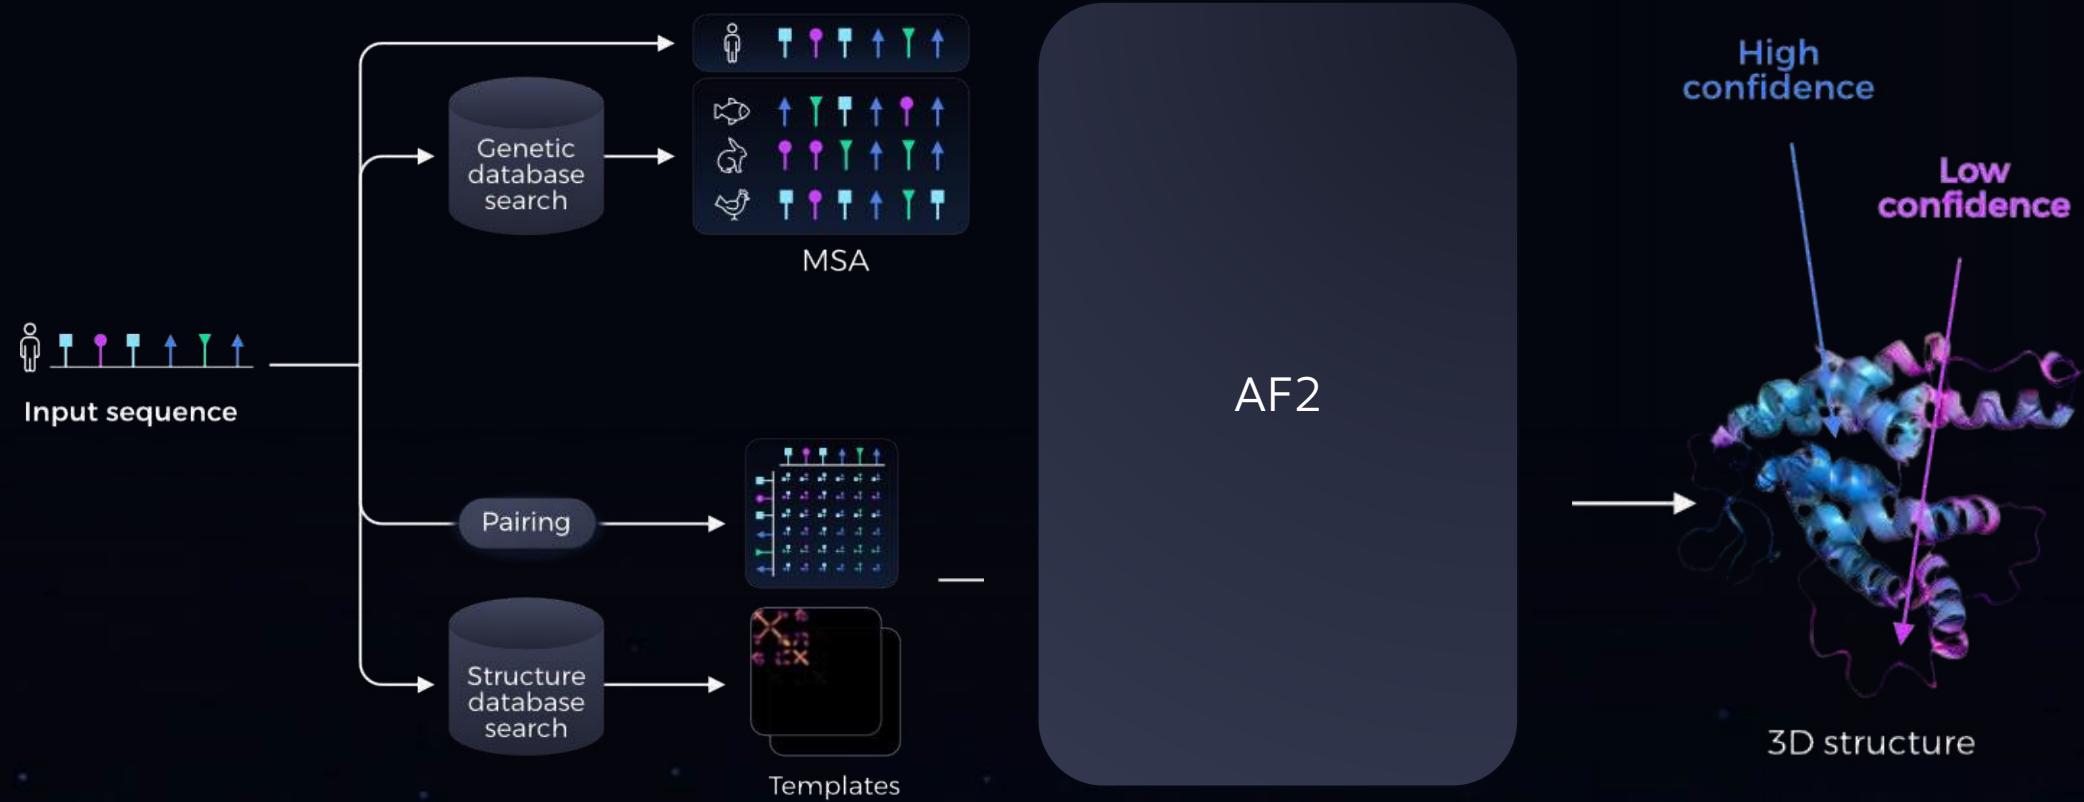


Source: Kovtun et al.  
PINDER, 2024 (VantAI)

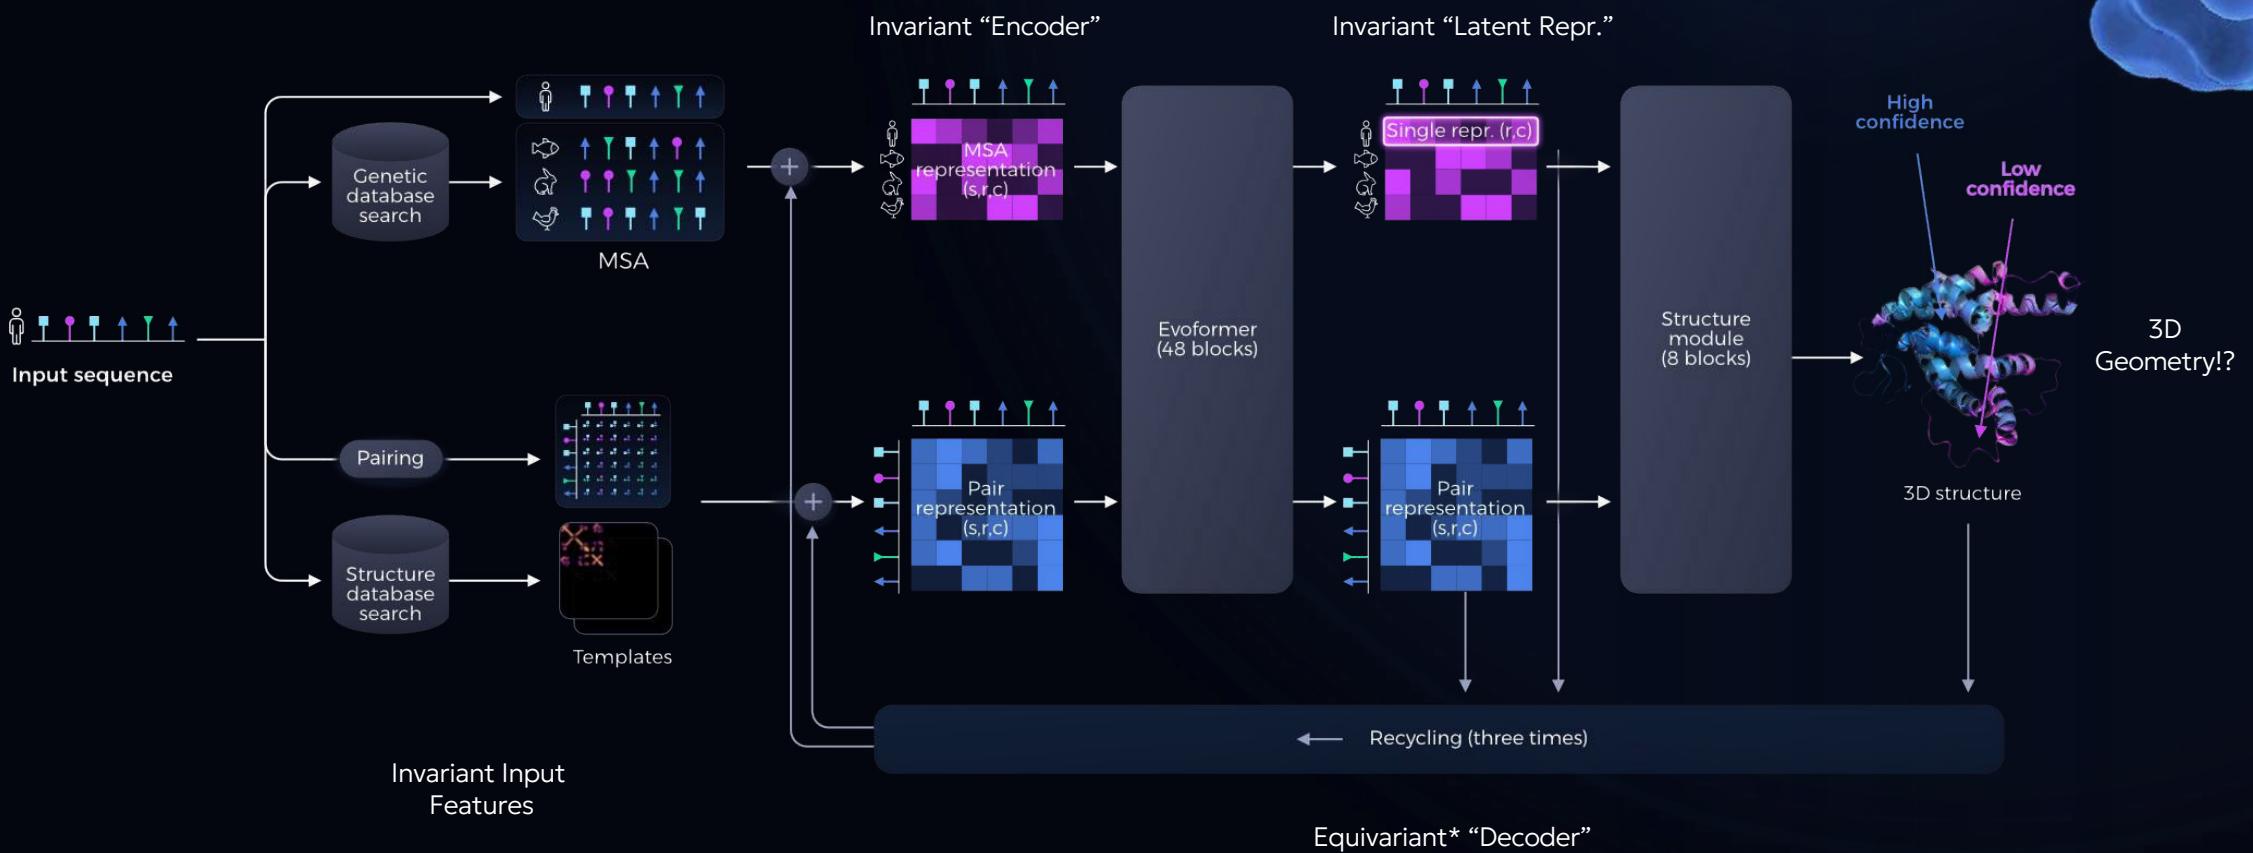
### Notes:

- For both monomers and protein-protein interfaces, MSA are absolutely critical information
- While current language models such as ESM3 have been shown to learn the pairwise residue covariance implicitly, despite 100B+ they still underperform a sequence alignment

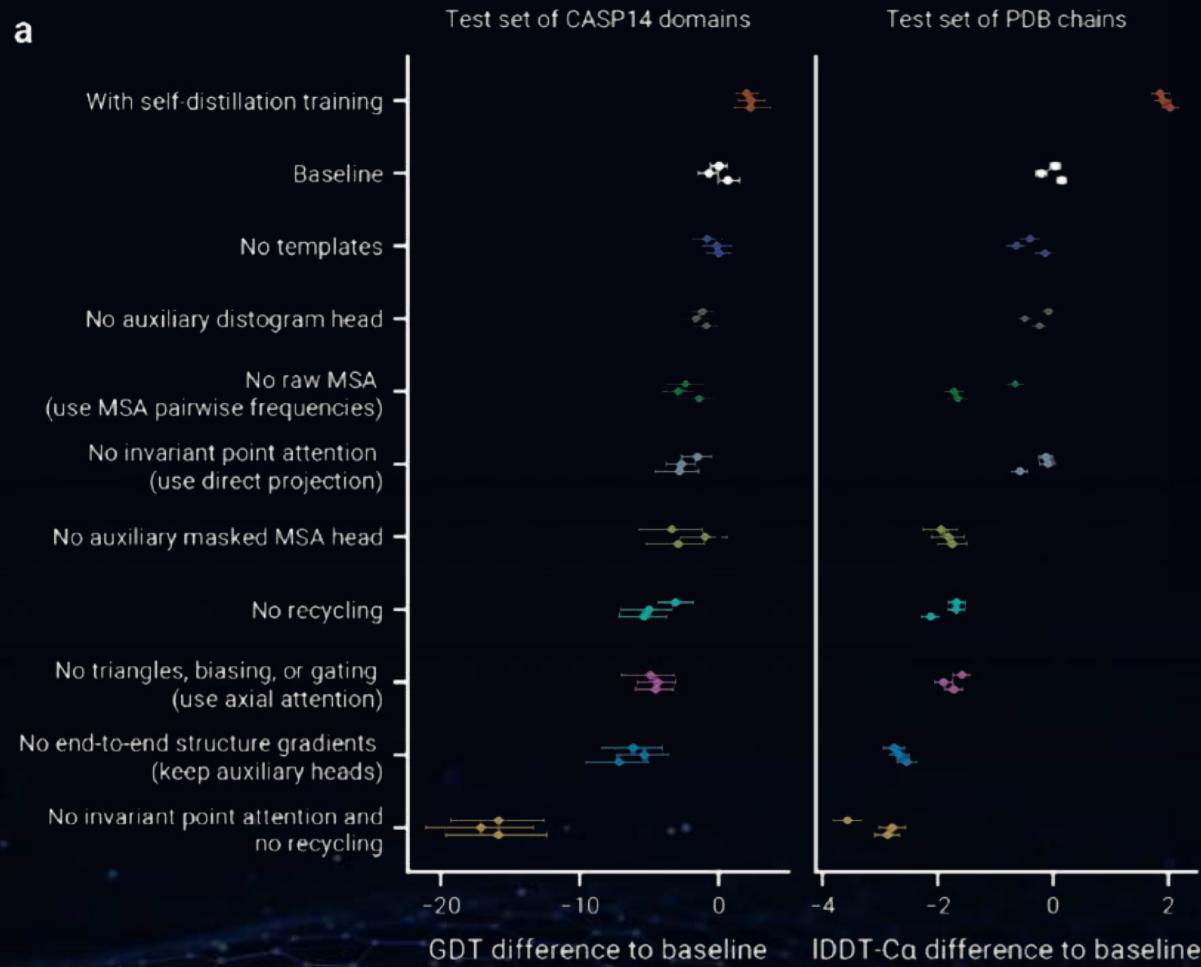
AF2 is a model that uses sequence and co-evolution to predict a protein's structure



It uses a transformer inspired architecture to first process sequence & MSA features and then reason over coordinates

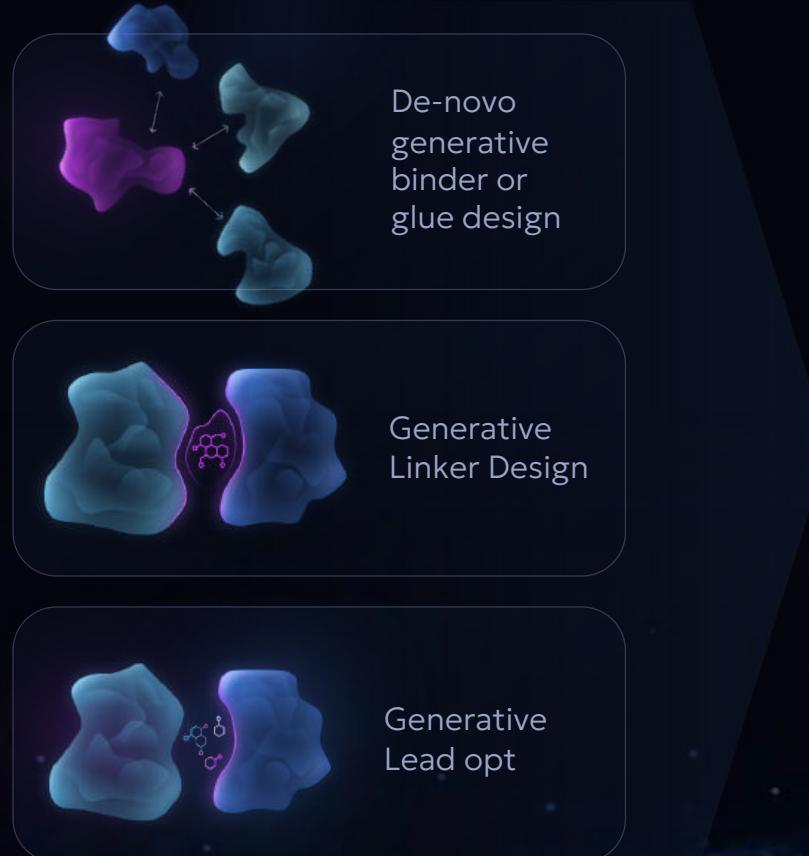


Critically – AF2 is an “engineering marvel” with many, many performance critical innovations



# De-novo molecular generators for variety of use cases

## Example use-cases for Proximity Modulators:



## Generative Algorithms



Diffusion/Flow Matching  
Probabilistic Models



LLMs +  
RL/RLHF



Genetic algorithms  
or search-based  
generative methods

## Scoring functions



Shape-based



Free Energy  
& interaction-based



ADME/PK  
Property classifiers

# MaSIF & dMaSIF have been highly successful for protein-design applications

## De novo design of protein interactions with learned surface fingerprints

<https://doi.org/10.1038/s41586-023-05993-x>

Received: 16 June 2022

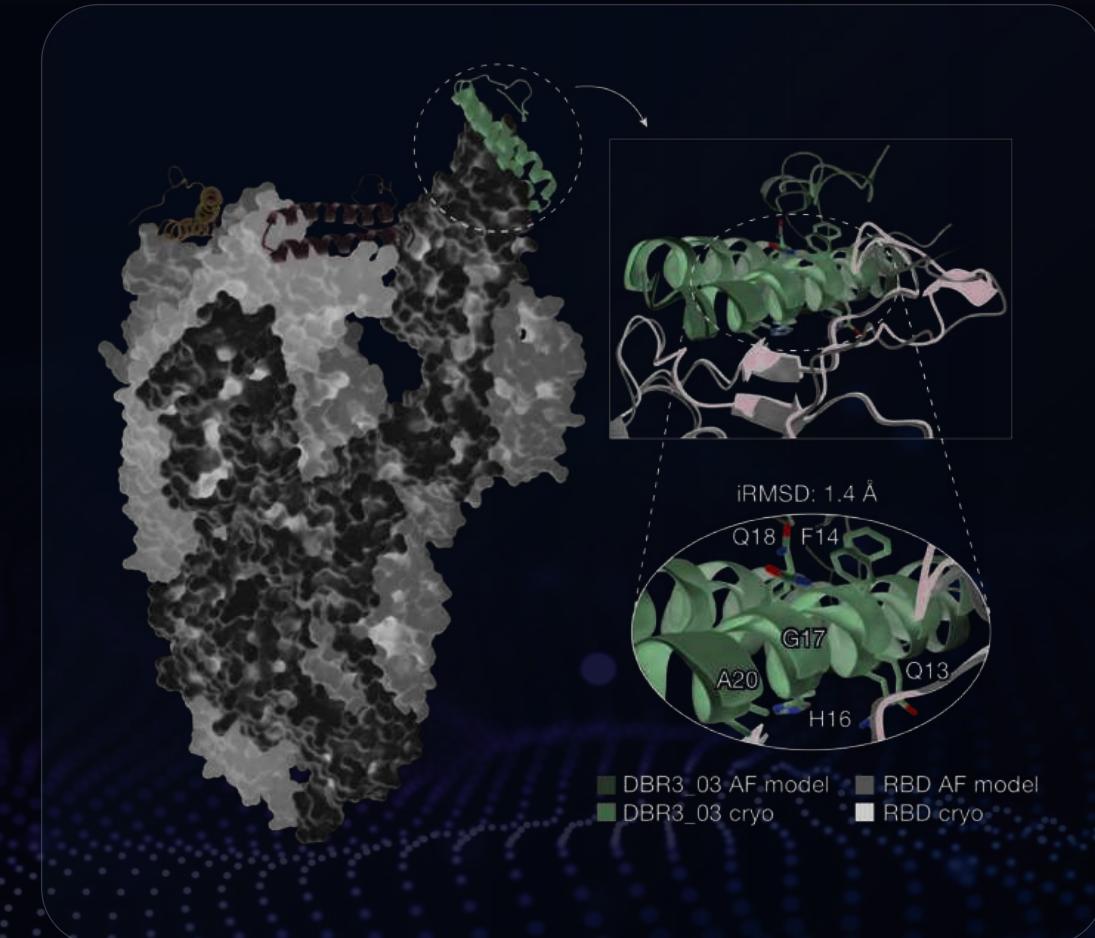
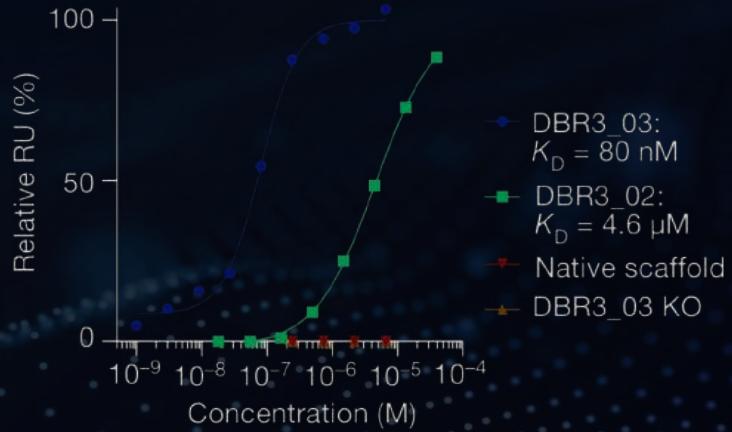
Accepted: 21 March 2023

Published online: 26 April 2023

Open access

Check for updates

Pablo Gainza<sup>1,2,10†</sup>, Sarah Wehrle<sup>1,2,11</sup>, Alexandra Van Hall-Beauvais<sup>1,2,11</sup>, Anthony Marchand<sup>1,2,11</sup>, Andreas Scheck<sup>1,2,11</sup>, Zander Harteveld<sup>1,2</sup>, Stephen Buckley<sup>1,2</sup>, Dongchun Ni<sup>3,4</sup>, Shuguang Tan<sup>5</sup>, Freyr Svärrisson<sup>1,2</sup>, Casper Goverde<sup>1,2</sup>, Priscilla Turrelli<sup>6</sup>, Charlène Racot<sup>6</sup>, Alexandra Teslenko<sup>1</sup>, Martin Pacesa<sup>1,2</sup>, Stéphane Rosset<sup>1,2</sup>, Sandrine Georgeon<sup>1,2</sup>, Jane Marsden<sup>1,2</sup>, Aaron Petruzzella<sup>5</sup>, Kefang Liu<sup>5</sup>, Zepeng Xu<sup>5</sup>, Yan Chai<sup>5</sup>, Pu Han<sup>5</sup>, George F. Gao<sup>5</sup>, Elisa Oricchio<sup>5</sup>, Beat Fierz<sup>7</sup>, Didier Trono<sup>5</sup>, Henning Stahlberg<sup>1,4</sup>, Michael Bronstein<sup>9,10</sup> & Bruno E. Correia<sup>1,2,11</sup>



A variety of generative algorithms have been used,  
usually split into 2D or 3D representations

### Generative Algorithms



Diffusion/Flow Matching  
Probabilistic Models

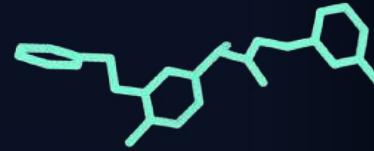


LLMs +  
RL/RLHF



Genetic algorithms  
or search-based  
generative methods

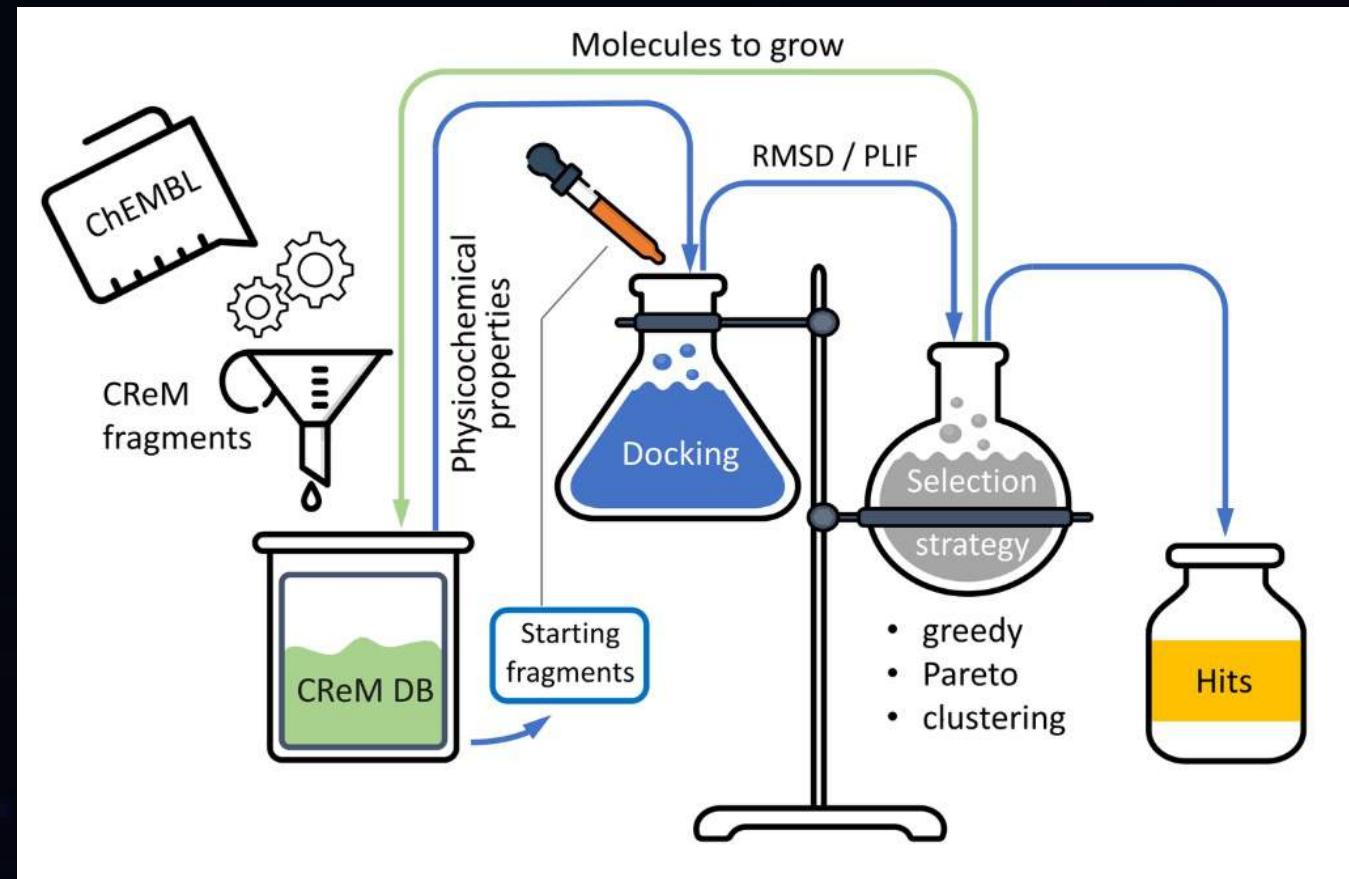
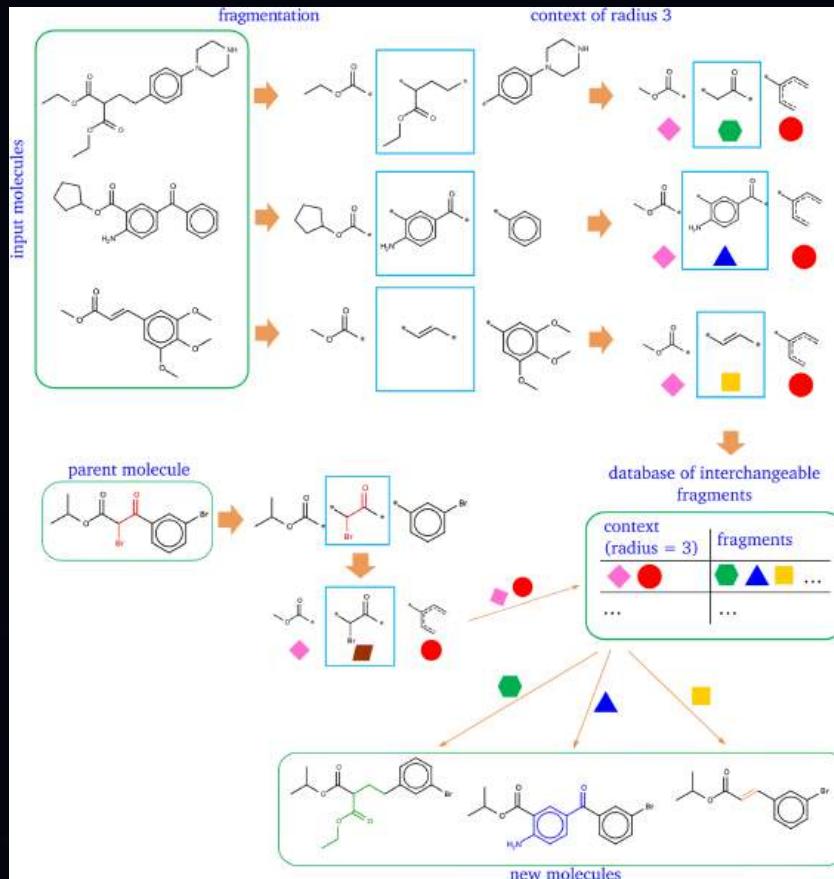
3D



2D

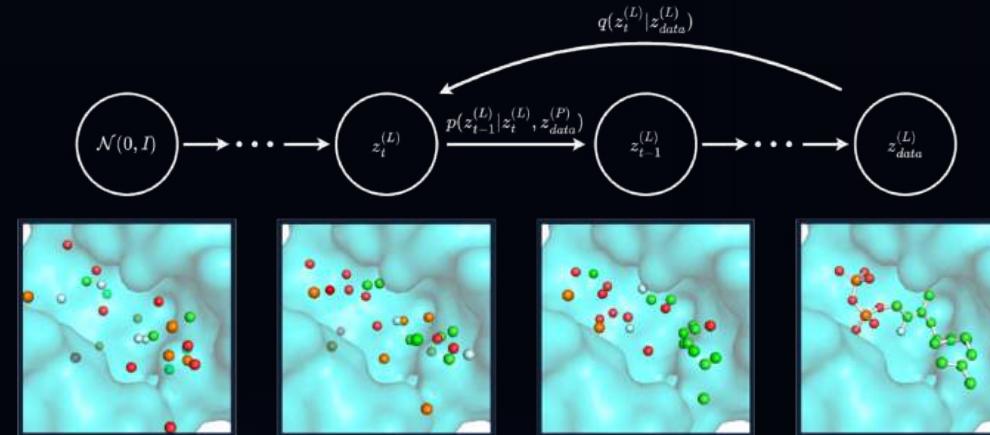
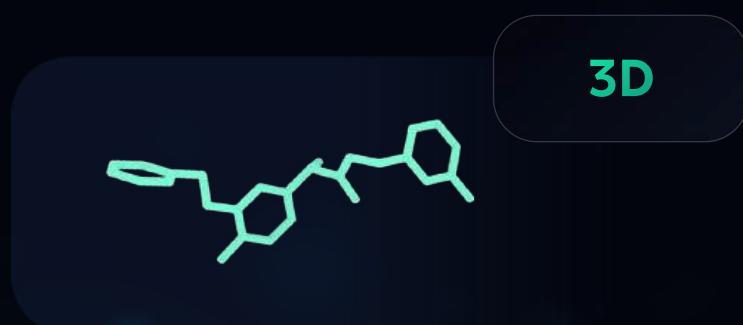
CC(O)CN

# Classical approaches: rule-based (iterative) enumerators

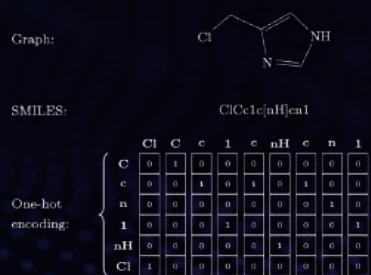


<https://github.com/ci-lab-cz/crem-dock>

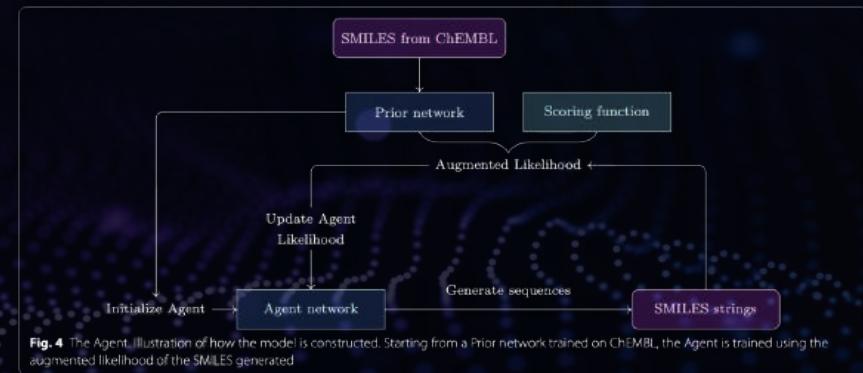
A variety of generative algorithms have been used, usually split into 2D or 3D representations



Structure-based Drug Design with Equivariant Diffusion Models. Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du, Carla Gomes, Tom Blundell, Pietro Lio, Max Welling, Michael Bronstein, Bruno Correia. ArXiV, 2022

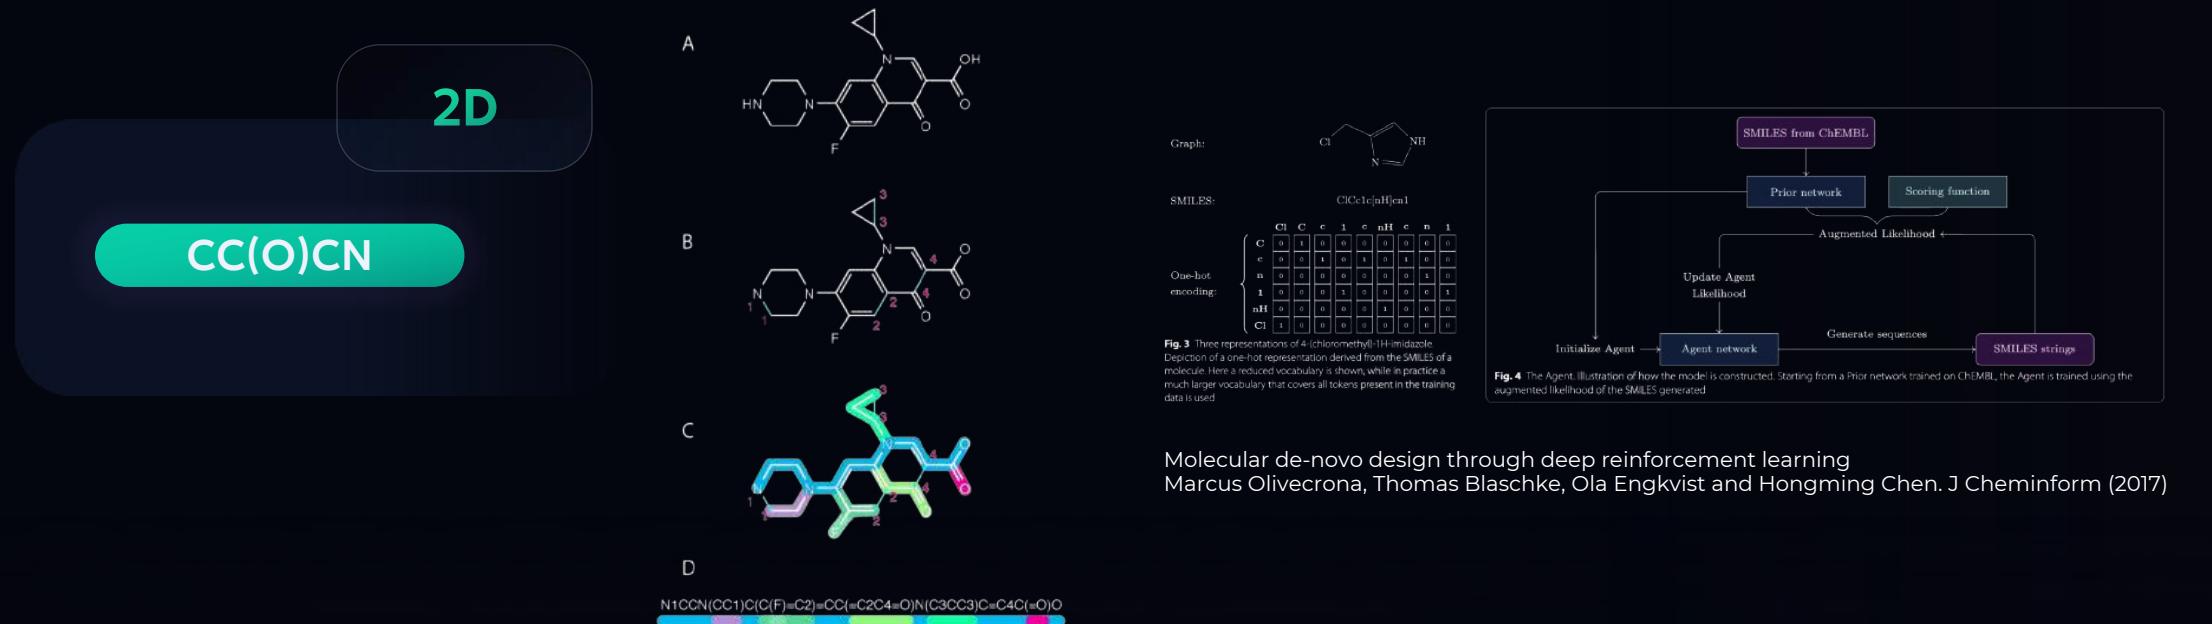


**Fig. 3** Three representations of 4-(chloromethyl)-1H-imidazole. Depiction of a one-hot representation derived from the SMILES of a molecule. Here a reduced vocabulary is shown, while in practice a much larger vocabulary that covers all tokens present in the training data is used.

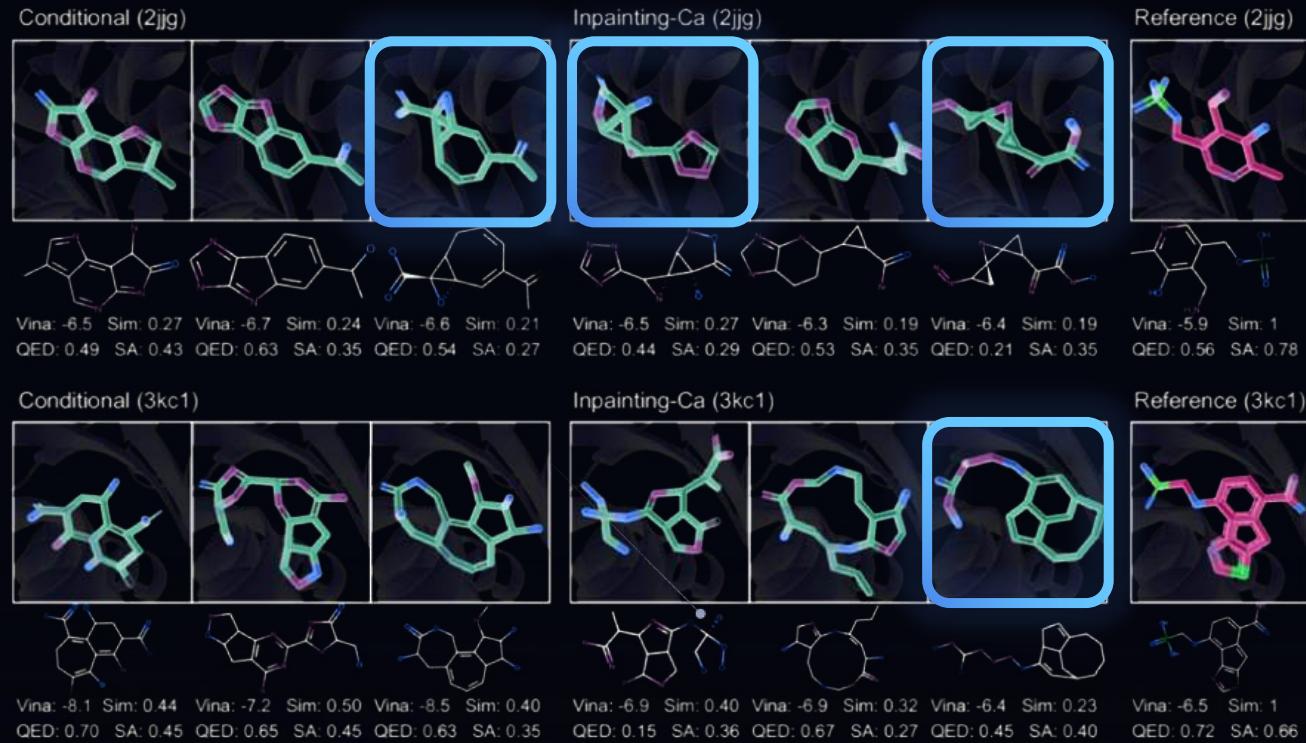


Molecular de-novo design through deep reinforcement learning  
Marcus Olivecrona, Thomas Blaschke, Ola Engkvist and Hongming Chen. J Cheminform (2017)

# SMILES is a string representation of a molecular graph which allows standard LLM tokenization approaches



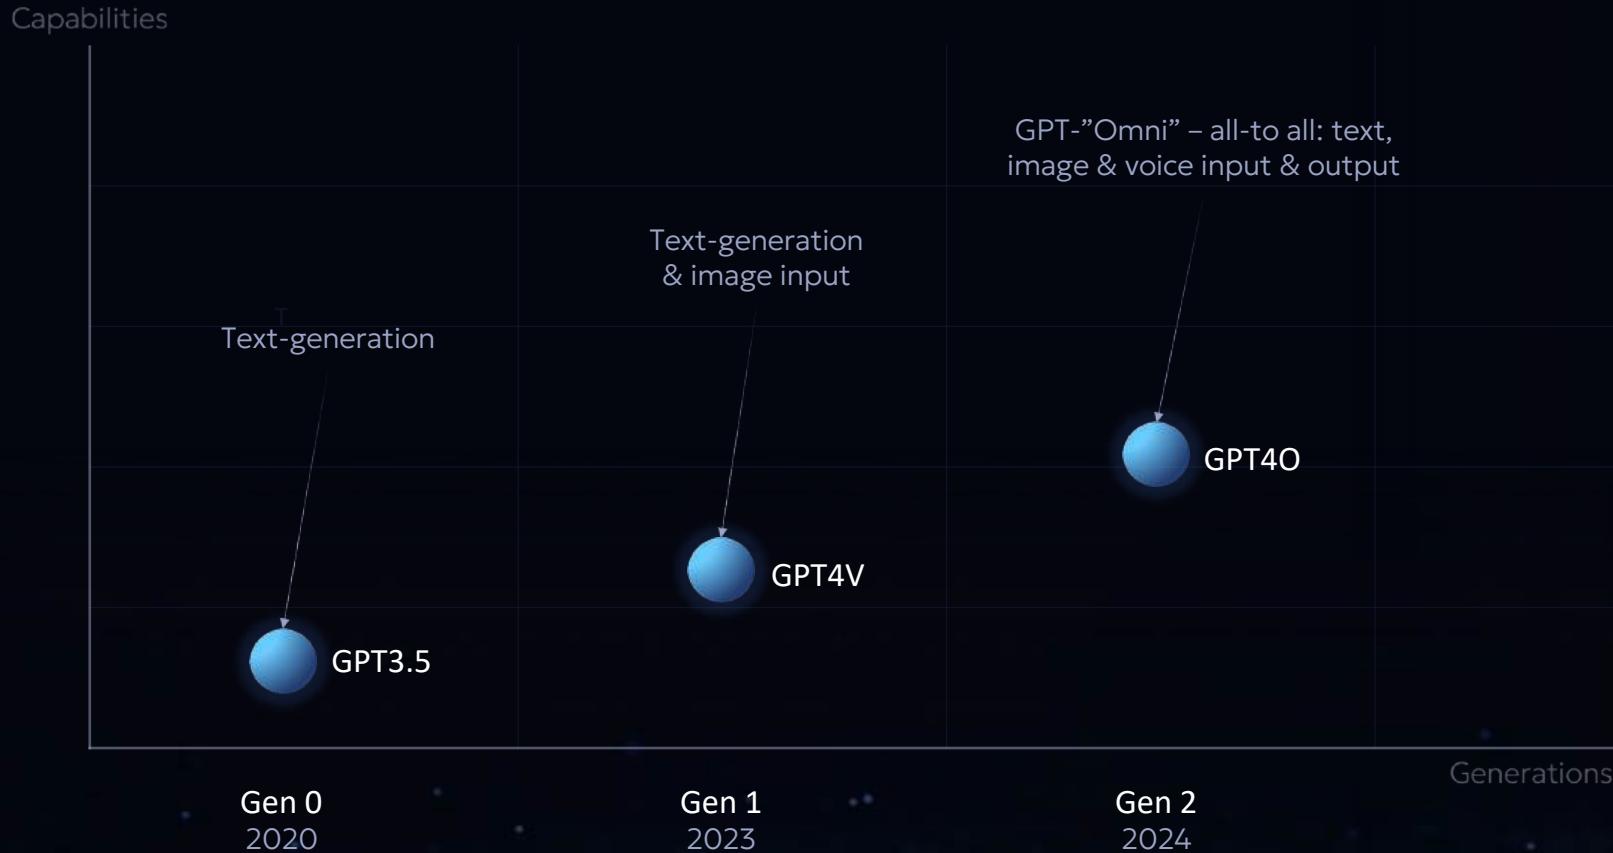
# 3D-based models still struggle to produce valid molecules



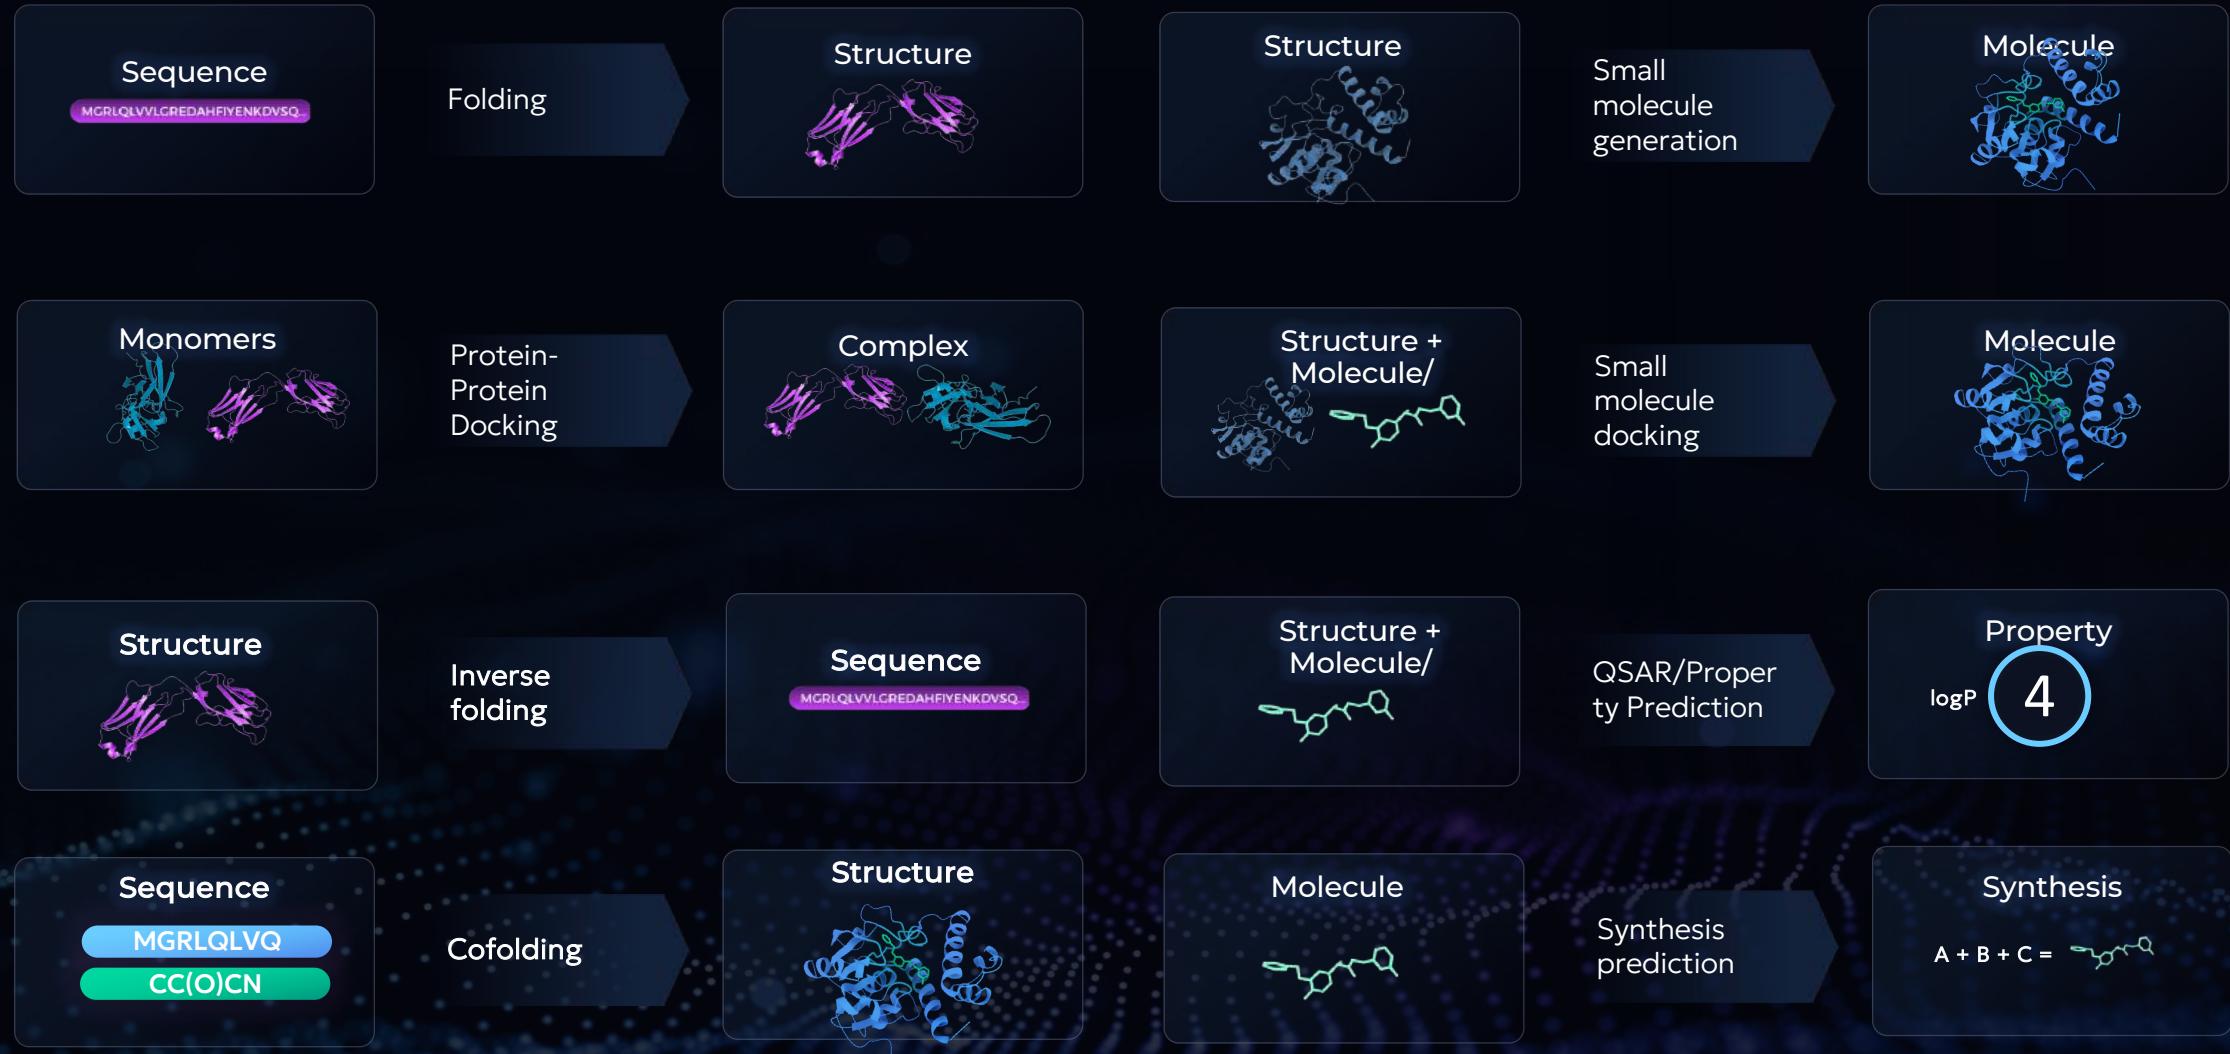
- Strained bond angles
- Highly strained or overextended rings
- Disconnected aromatic bonds
- Large bent overall molecular shapes
- Many methods are still in the “6-Finger phase” of generative models
- Whether same recipe from image world (larger models, more data: scale) will suffice or if explicit inductive biases will be required remains to be seen



The big breakthroughs in AI have been driven through ever more general models



Current Bio AI models, in contrast, still consist of a large number of specialist models



# However, this is wasteful and limiting

- Models re-learn same things anew ("grammar" of natural sequence, physics of interactions, generative process, ...)
- Different and small datasets used to re-learn across tasks
- In real life, we often want to "feed" models with whatever information we have
- Currently, this requires models to be chained or "hacked" to include information they are not trained to use (e.g. MSA sampling in AF2), creating and compounding errors

With AF3, all atoms of life (DNA, Protein, Molecules) could be predicted for the first time

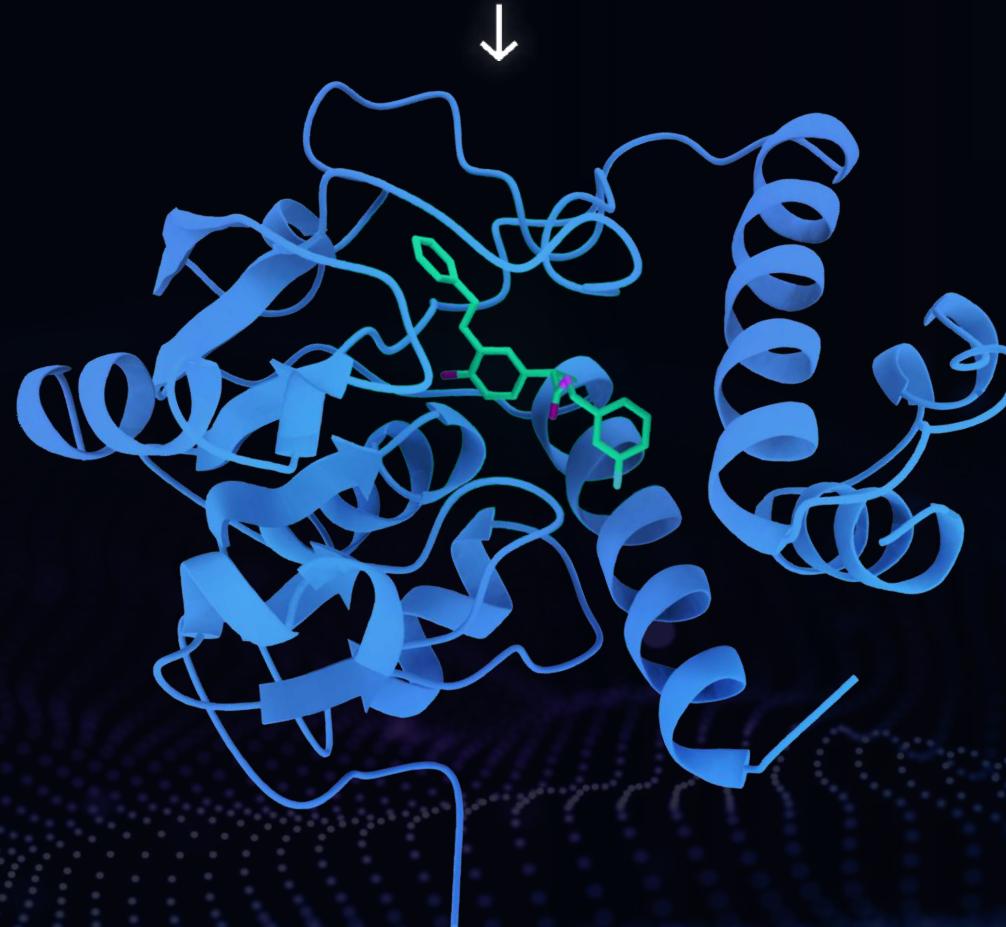


# Co-Folding:

Predict atomistic  
structure given  
sequence of  
molecular tokens

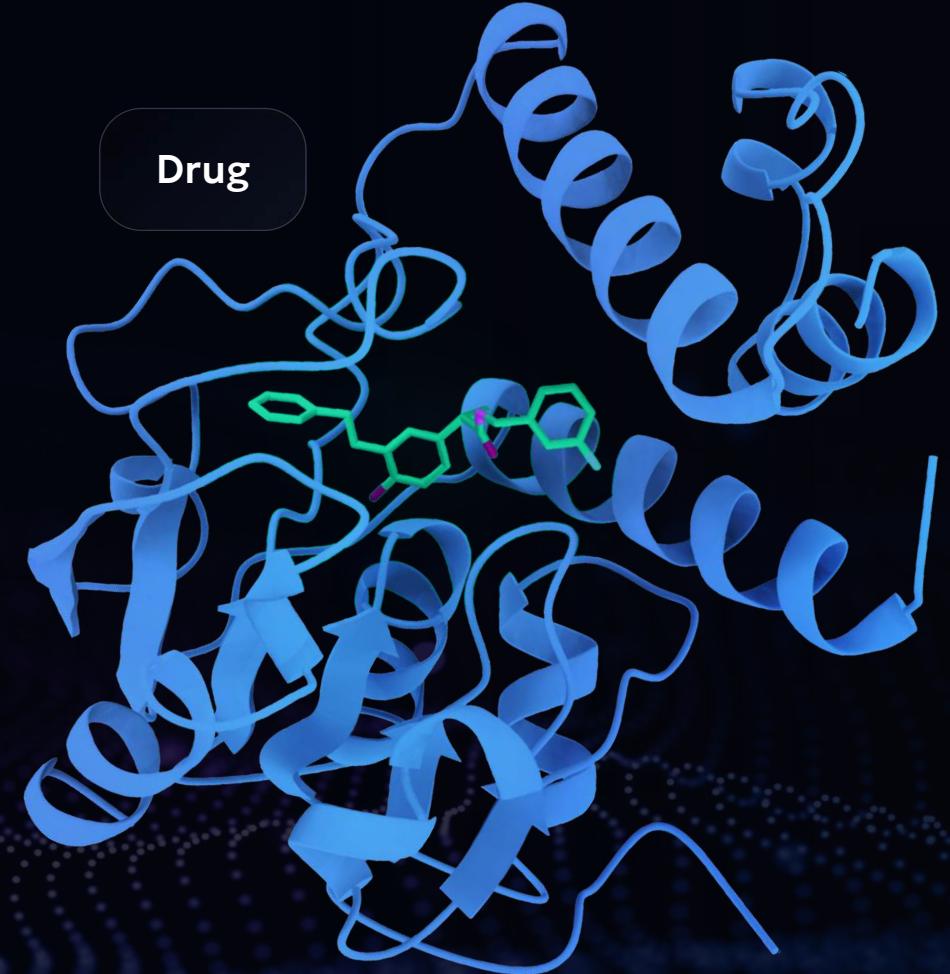
MGRLQLVQ

CC(O)CN



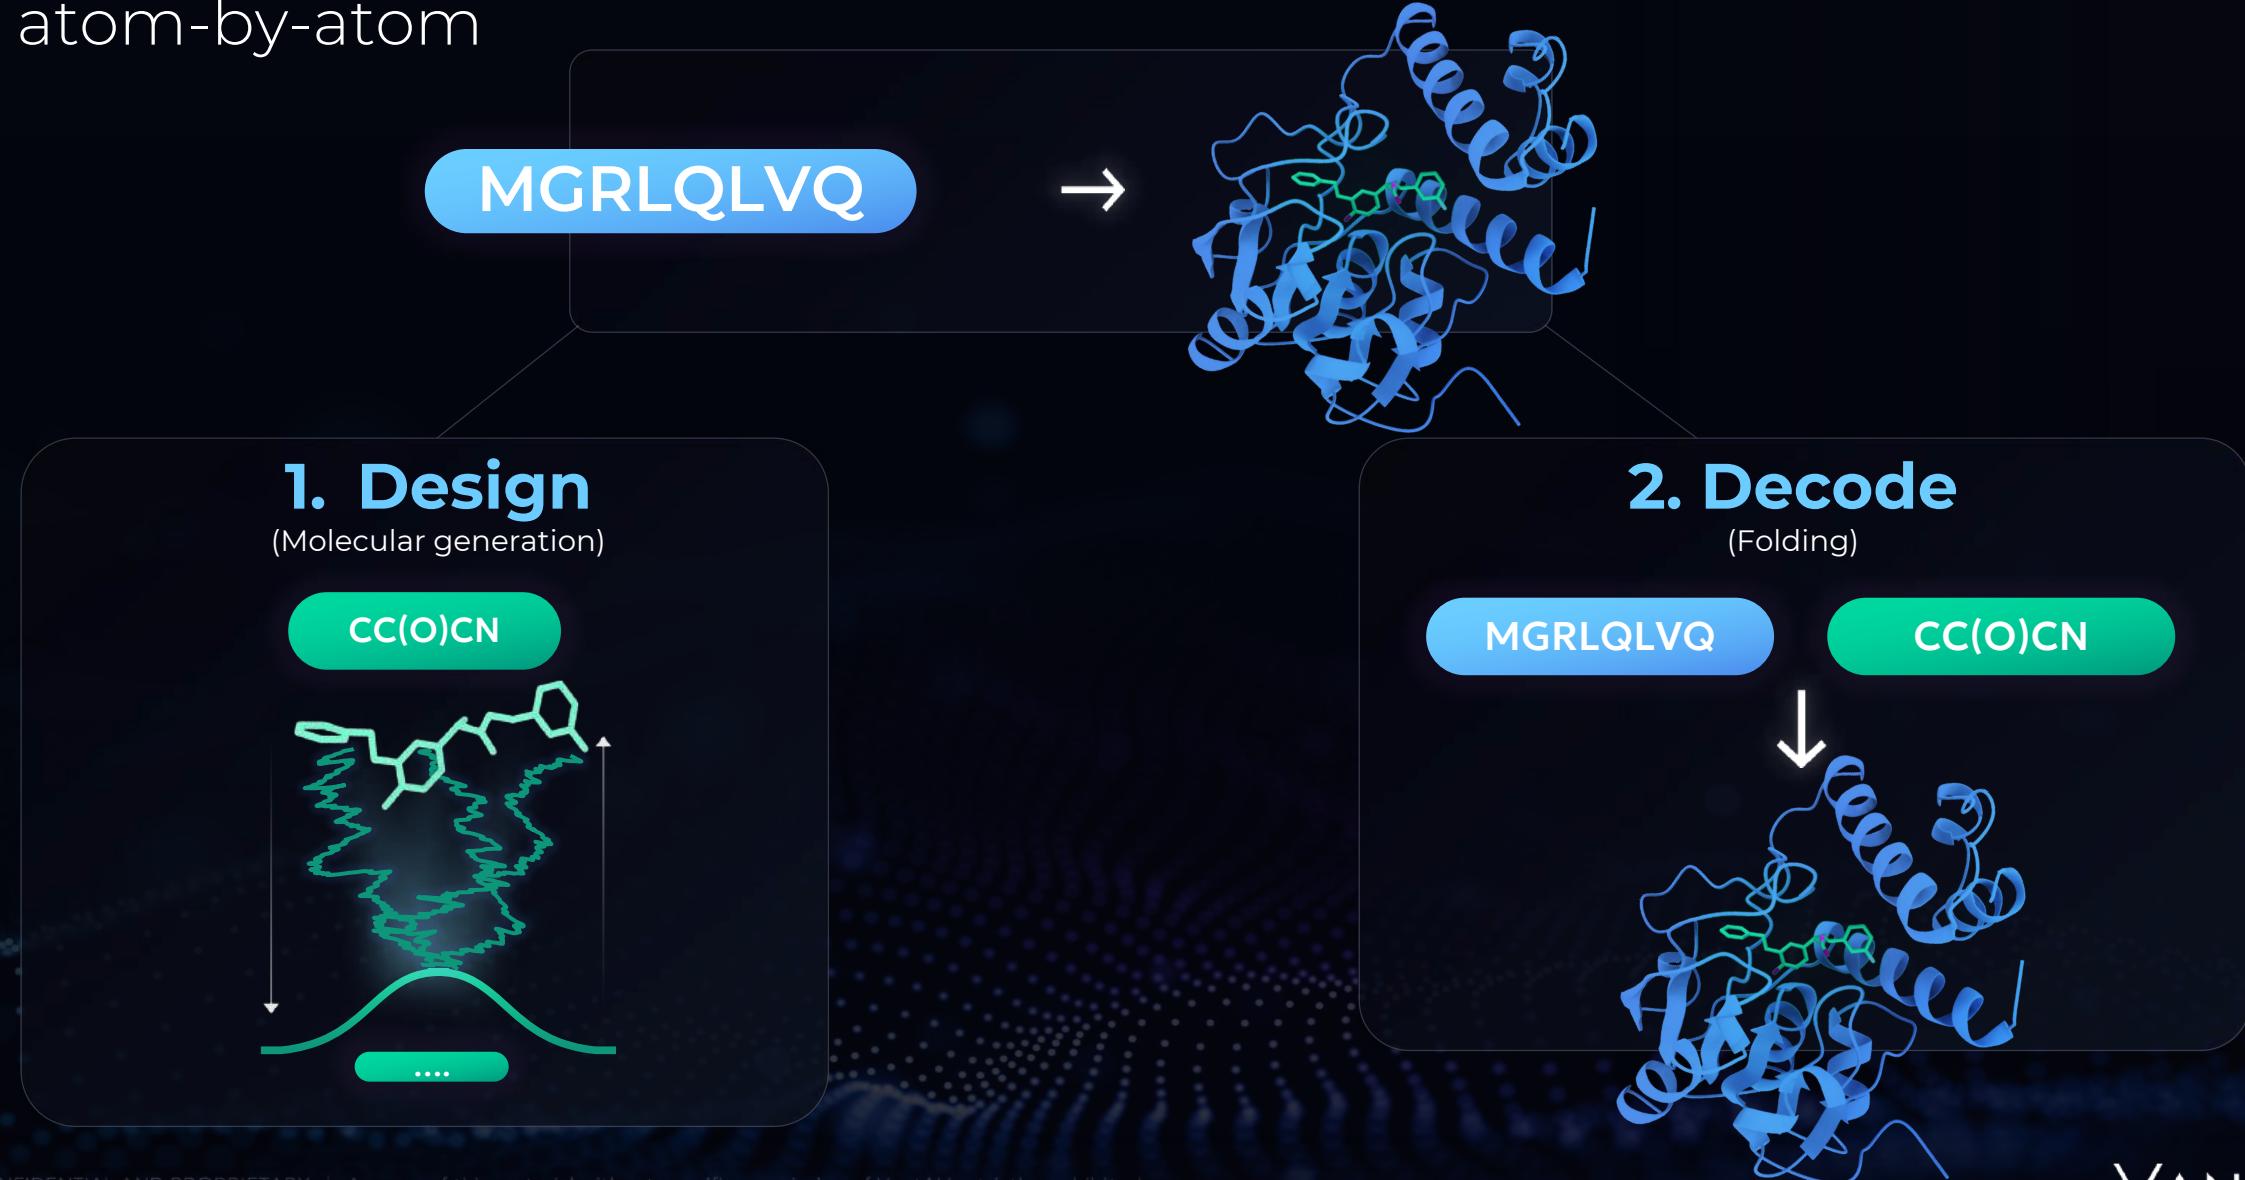
# The challenge Neo set out to solve:

**Decade+** process of finding an effective molecule



# Challenge: Design & decode medicines

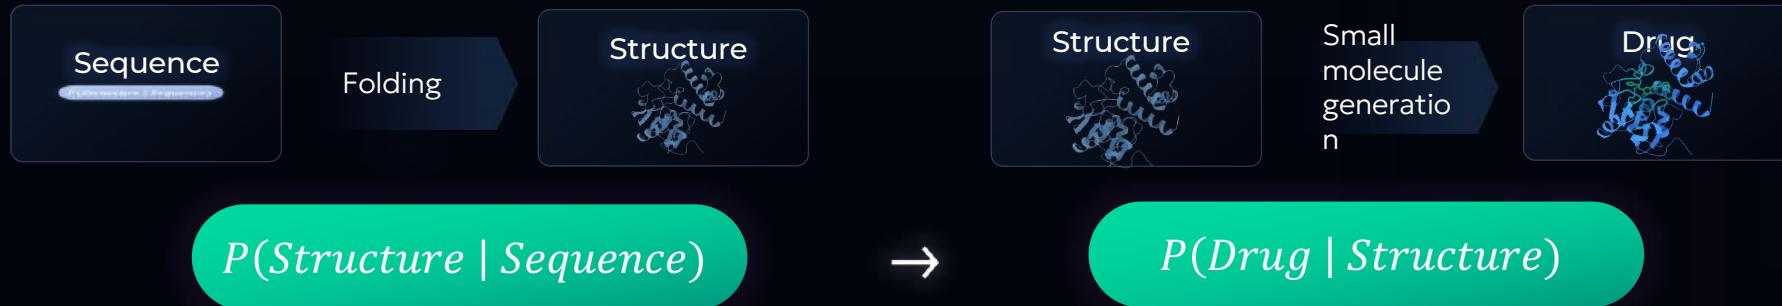
atom-by-atom



And the next frontier was clear



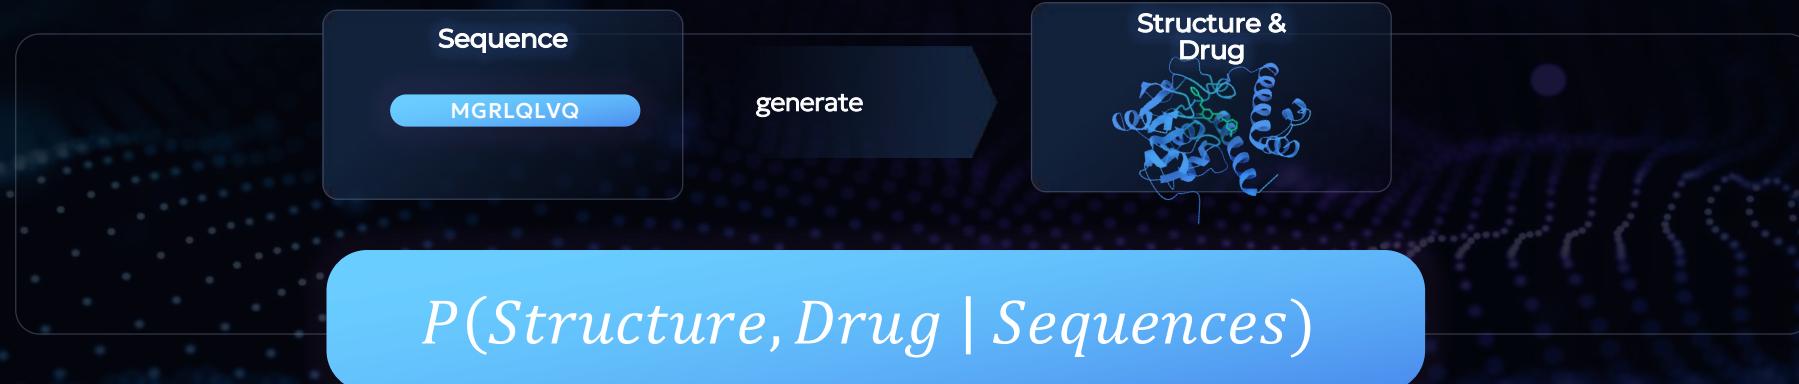
## Existing methods



However, in ProMods

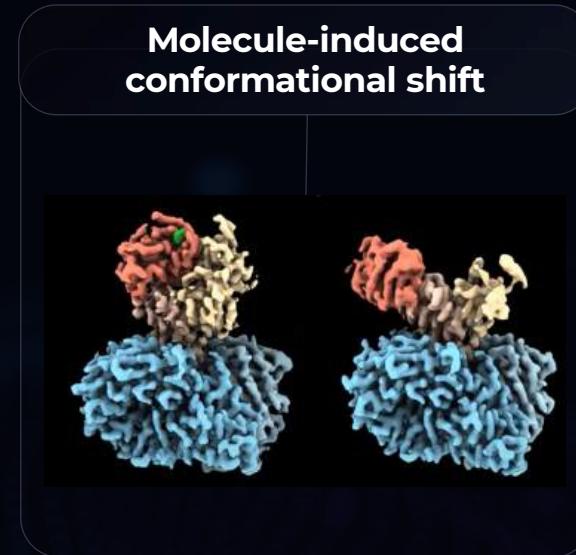
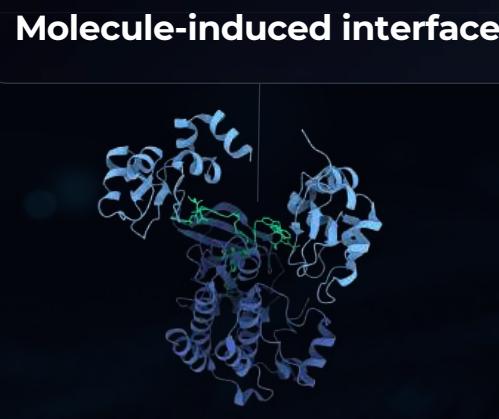
$$\text{Structure} = f(\text{Drug}, \text{Sequences})$$

## Needed

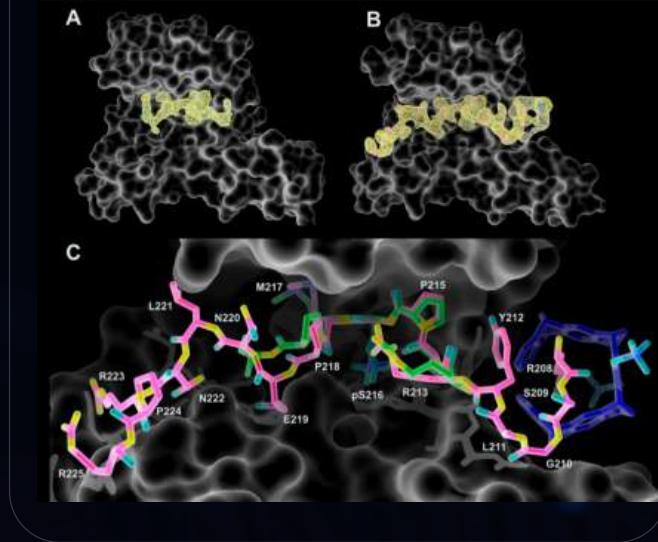


# ProMods: molecule defines structure

often **doesn't form** stably **in absence of drug**



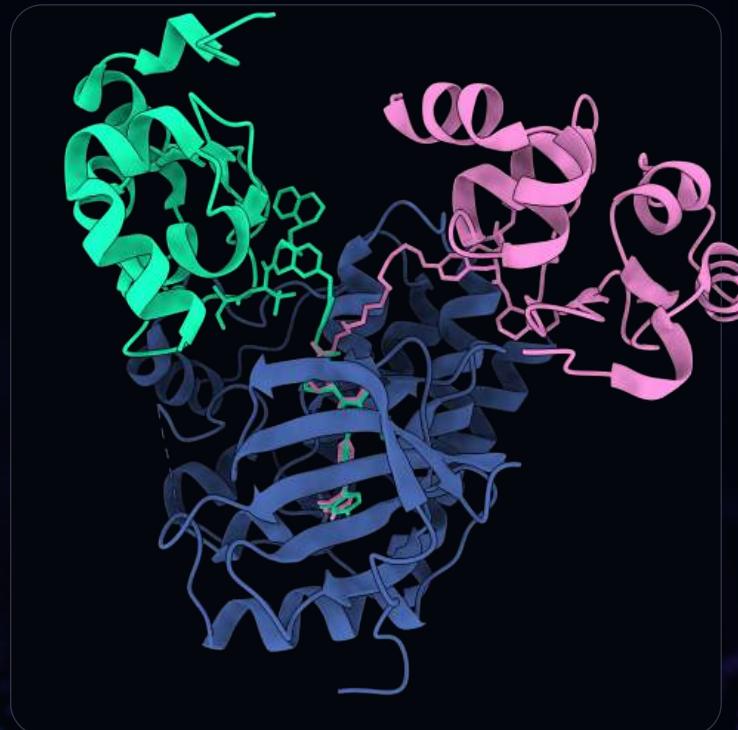
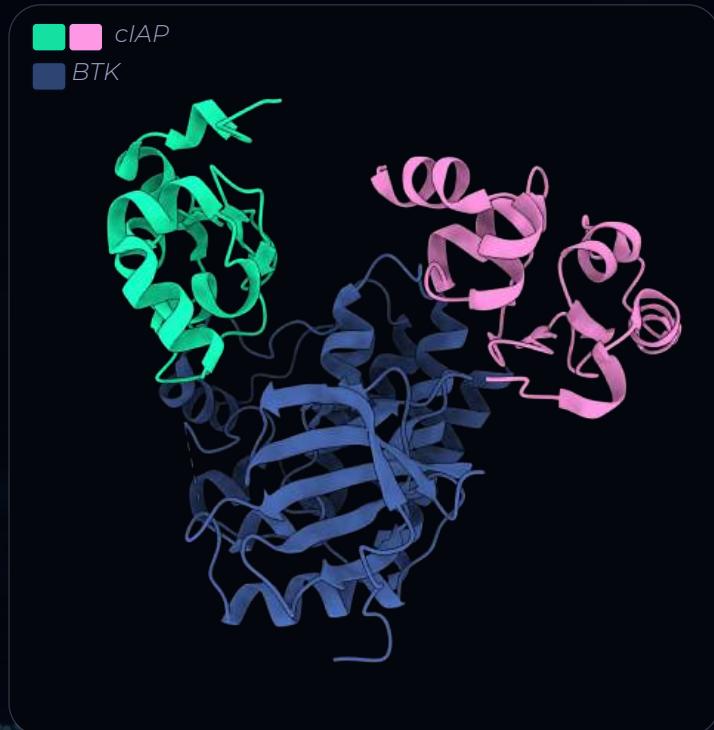
Molecule-induced disorder to order conversion



Protein-Protein & Protein-Ligand docking invalidly assume interface and monomers exist stably without molecule – simultaneous co-folding & de-novo design required

$$\text{Structure} = f(\text{Drug}, \text{Sequences})$$

Which of these two cIAP-BTK Protein-interfaces  
is the right one?

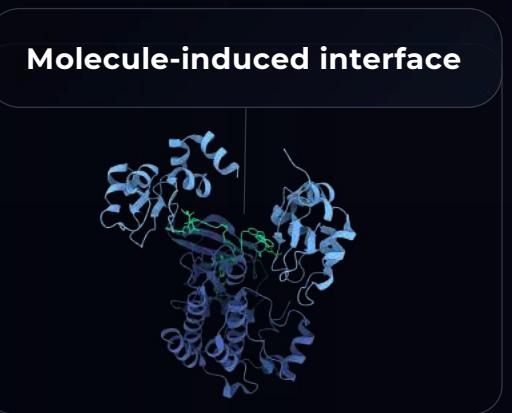
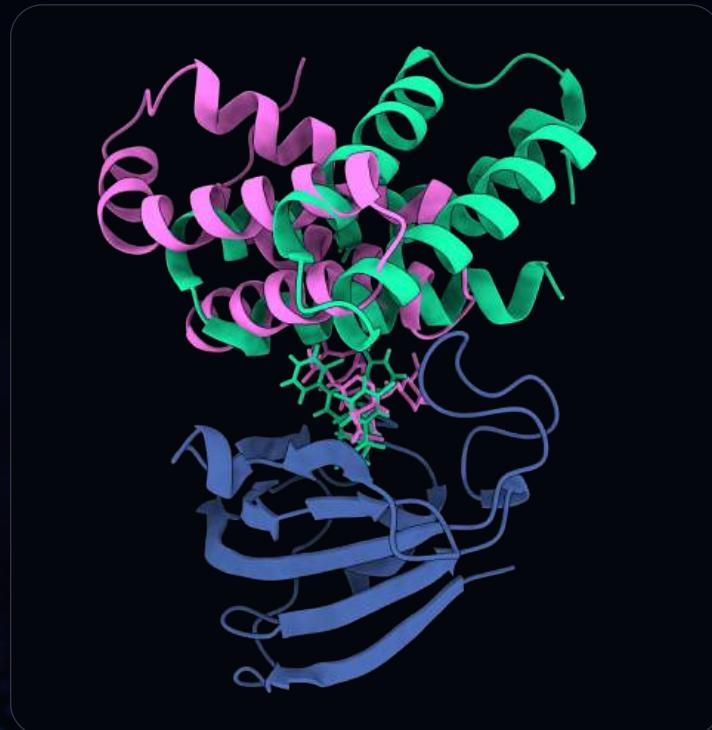


Both! Depending on molecule

**Molecule-induced interface**



Which of these two FKBP12-mTOR Protein-interfaces is the right one?



# De-novo design:

De-novo small molecule design still lag behind:  
a) methods assume knowledge of bound-state  
protein & b) struggle with valid designs

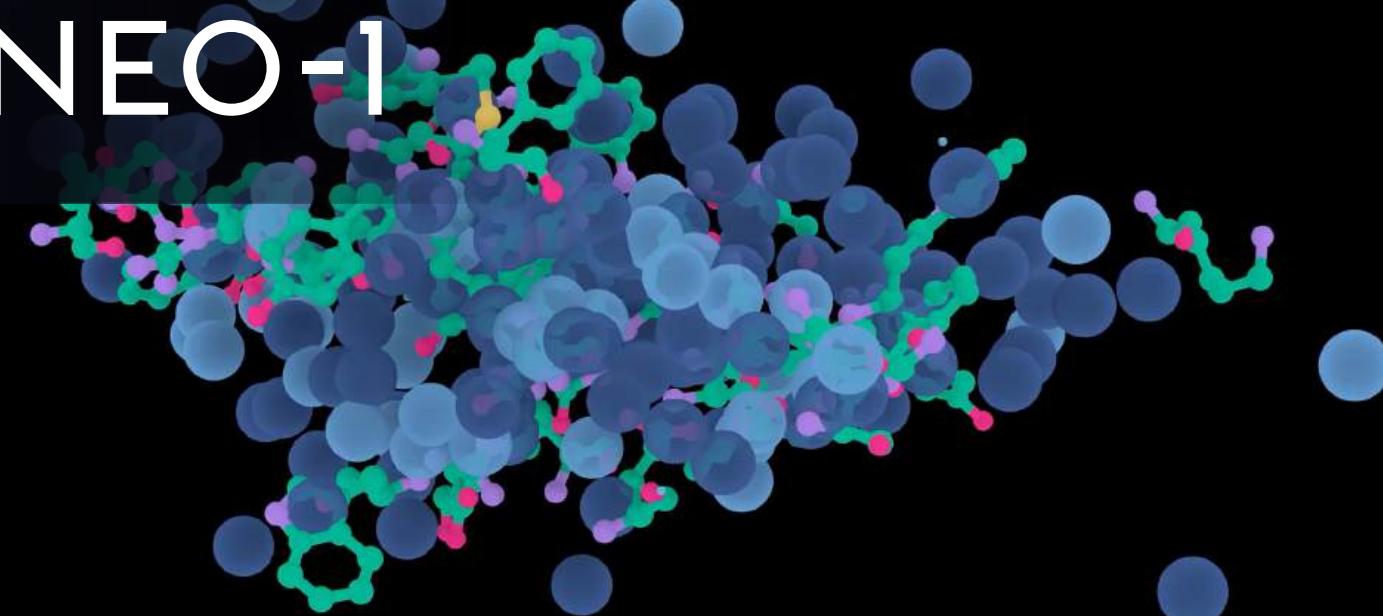


Structure-based Drug Design with Equivariant Diffusion  
Models. Schneuing, ..., Bronstein, Correia. 2022

Introducing



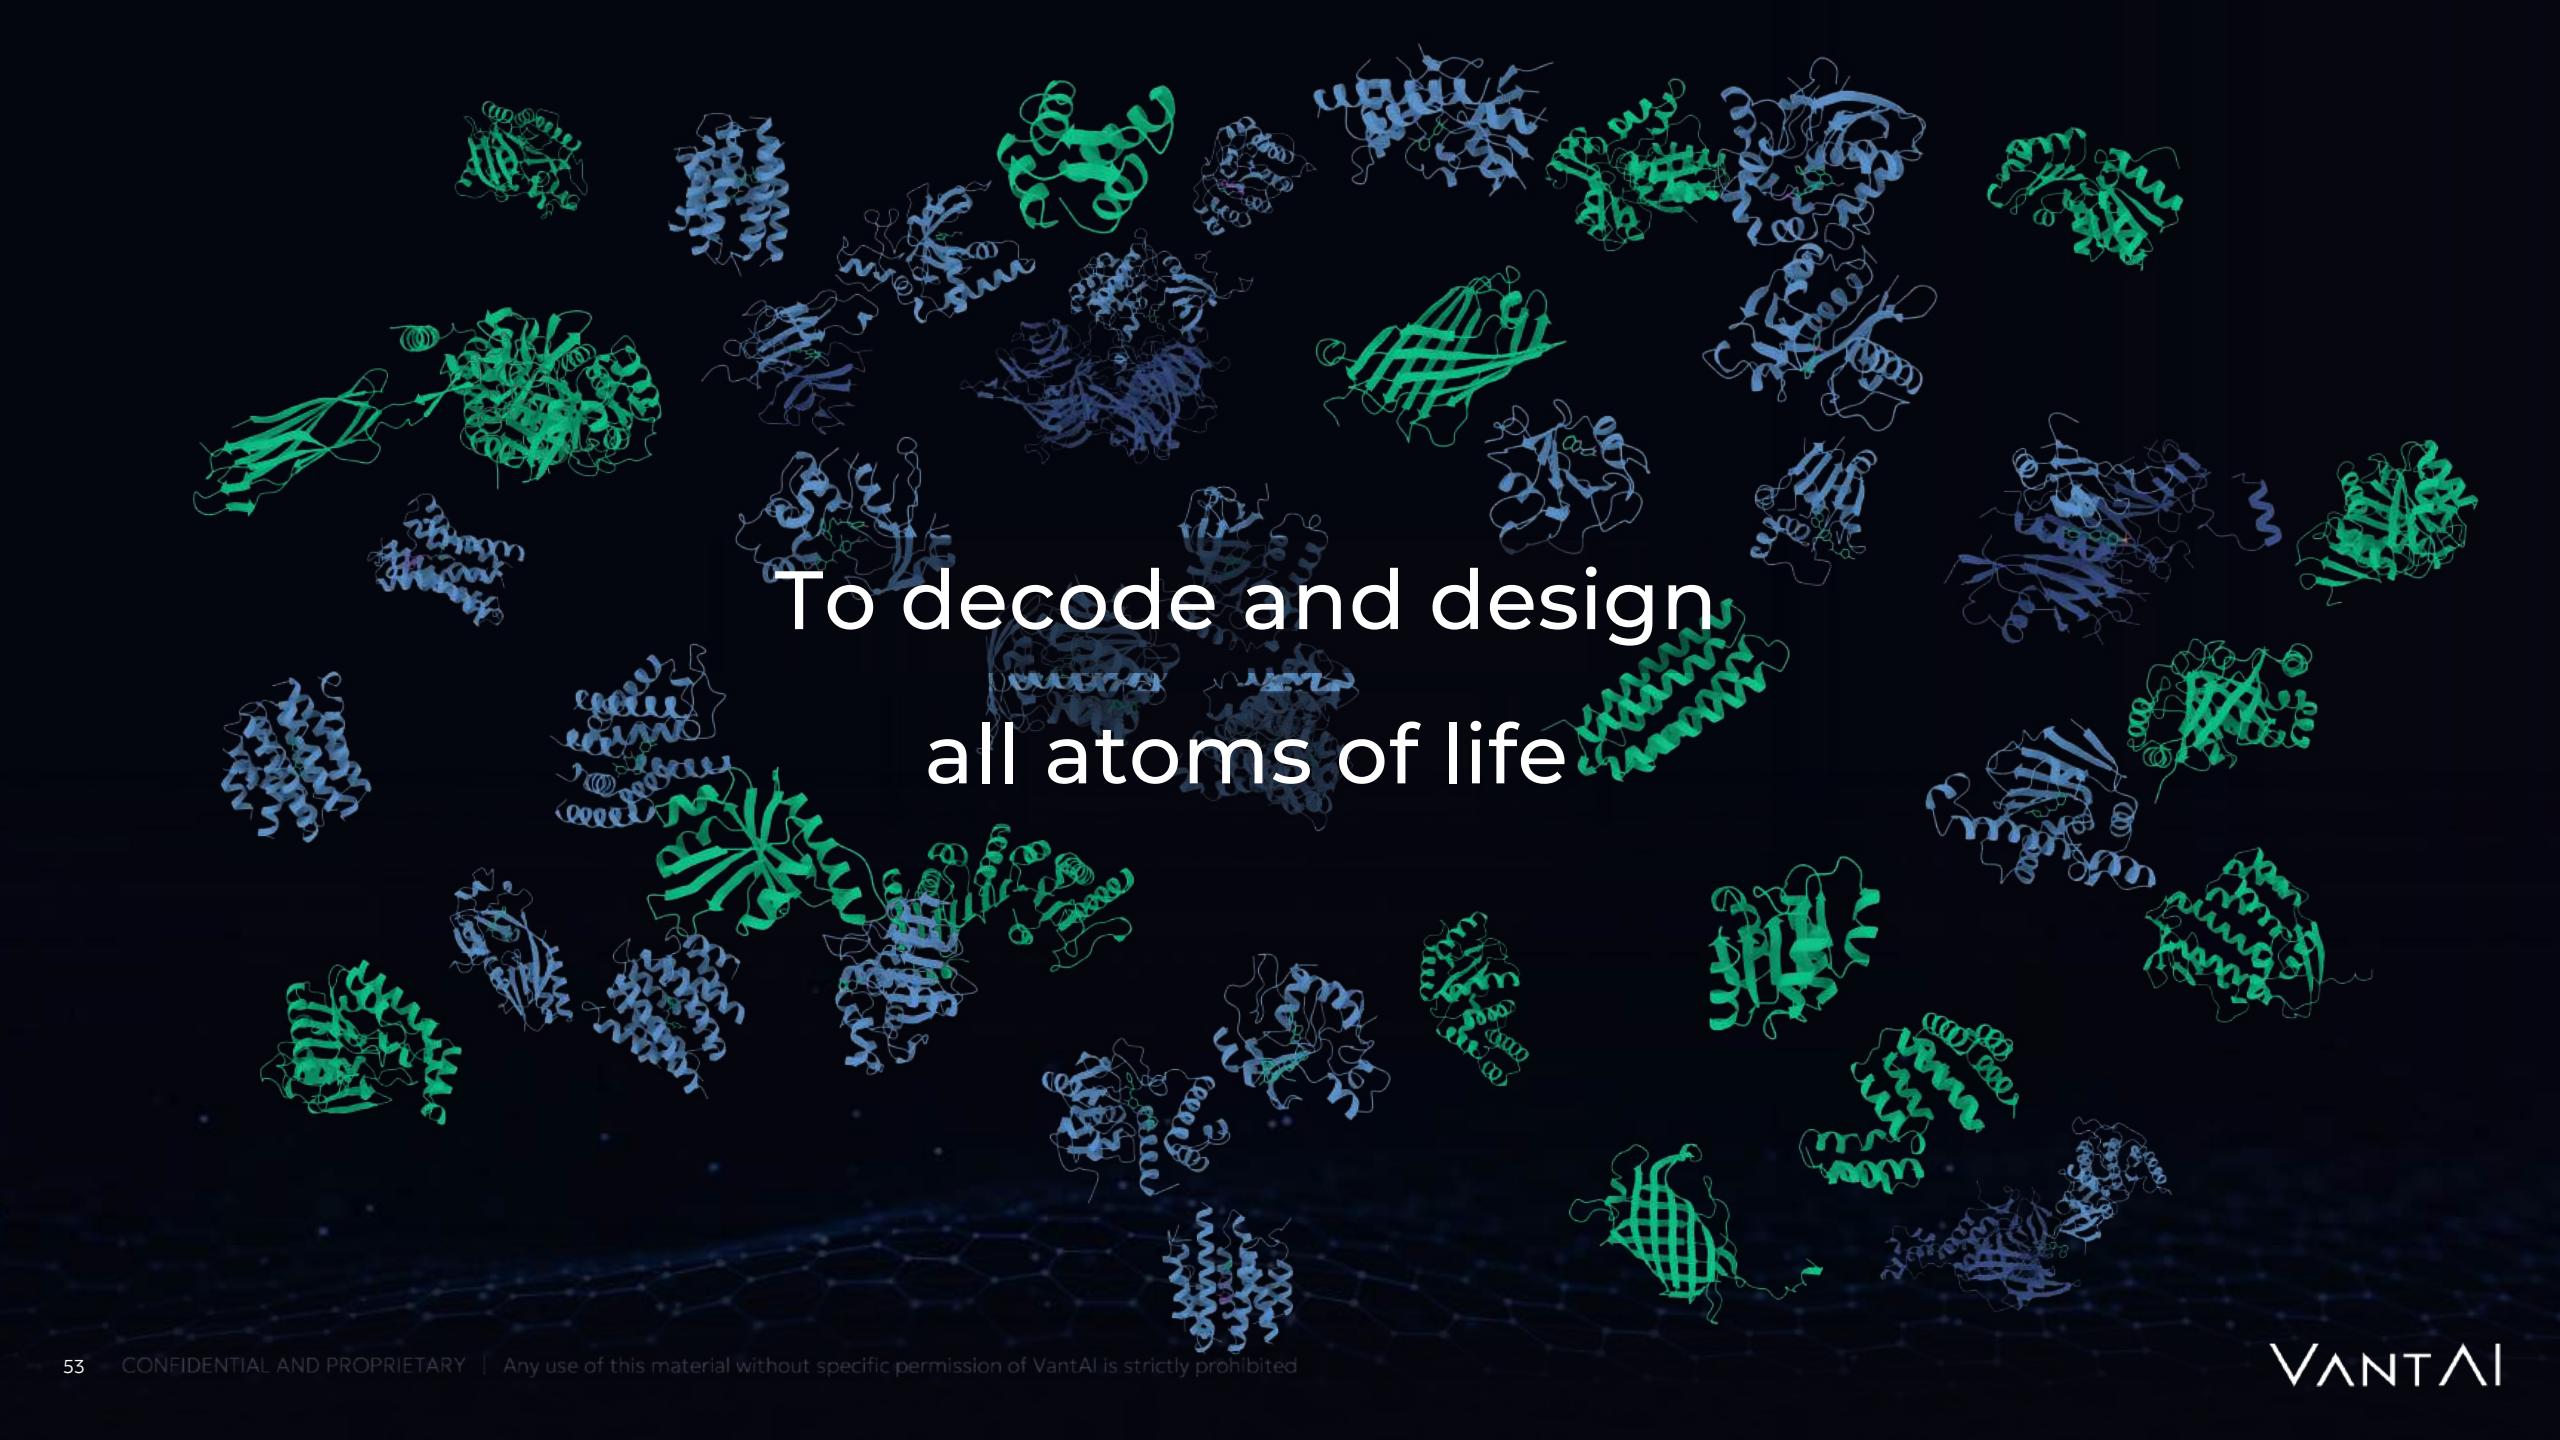
VANTAI  
NEO-1



[vant.ai/neo-1](https://vant.ai/neo-1)

# Introducing Neo-1: the worlds most advanced atomistic foundation model





To decode and design  
all atoms of life