


UNIVERSITÄT
HEIDELBERG



UNIVERSITY OF
CAMBRIDGE

Drug Design – Present and Future

L10, Structural Bioinformatics

WiSe 2023/24, Heidelberg University

Overview for this lecture


- 1. Background: Drug Discovery Pipeline**
- 2. Traditional approaches to early-stage design**
- 3. Deep learning-based docking methods**
- 4. Generative modelling for drug design**

1. Background: Drug Discovery

Eroom's Law


Drug Discovery is hard

a Overall trend in R&D efficiency (inflation-adjusted)



Recommended Reading:
- Derek Lowe – [‘In the pipeline’](#)

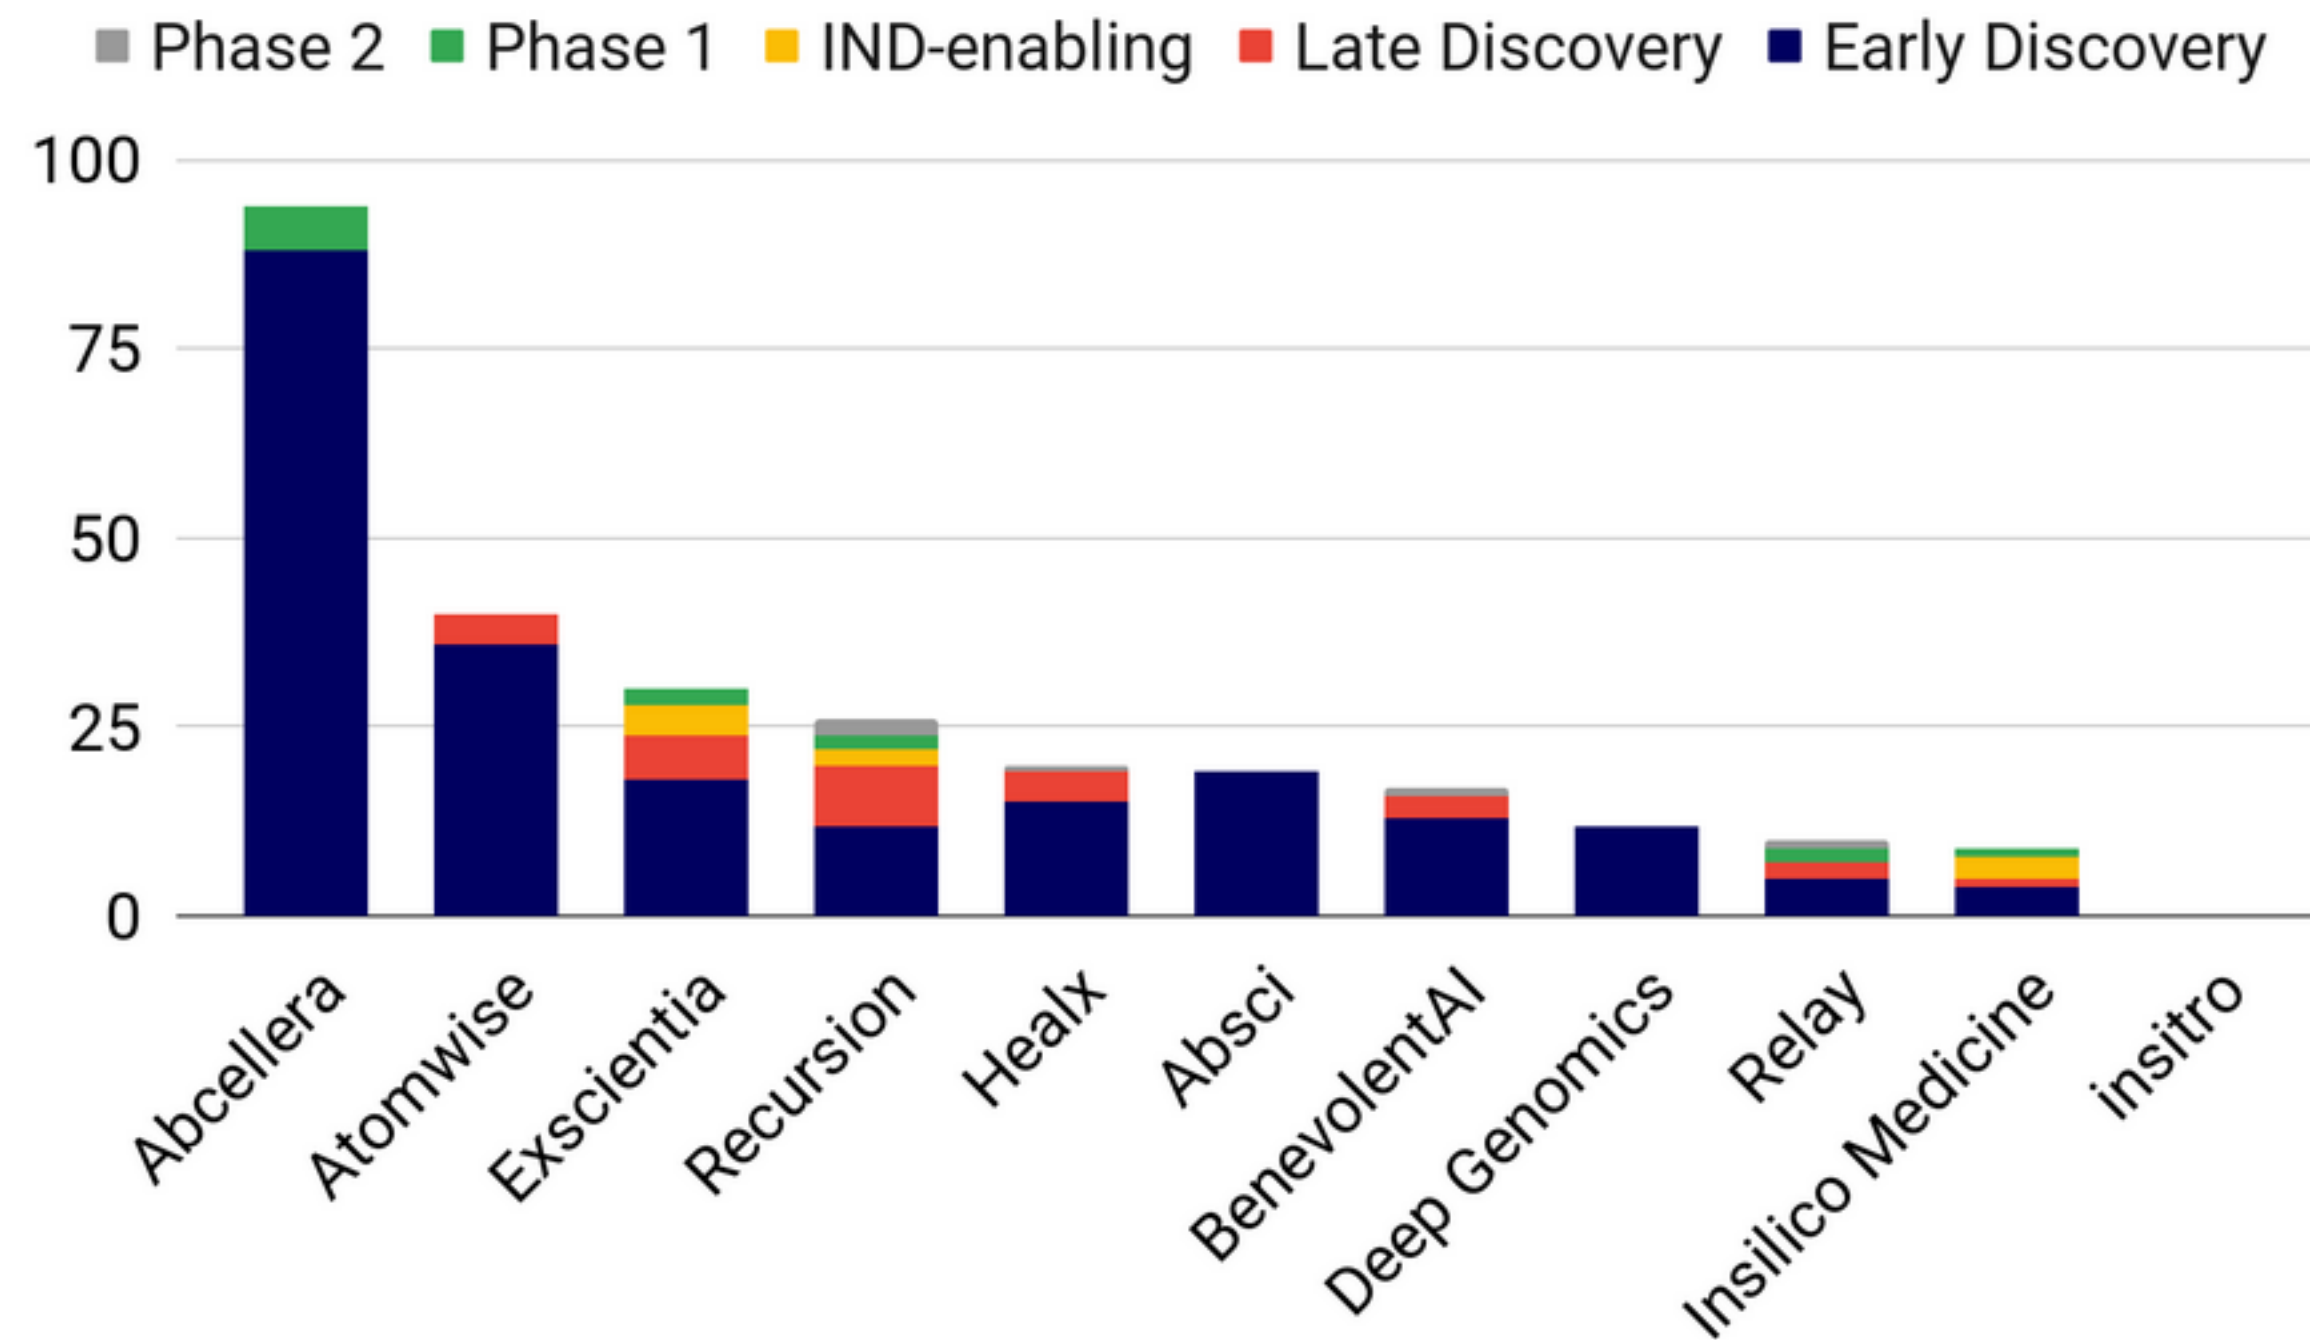
The Drug Discovery Pipeline




AI-first Drug Design – reducing the cost?

>18 assets from AI DD companies now in trails

of assets per pipeline stage per company



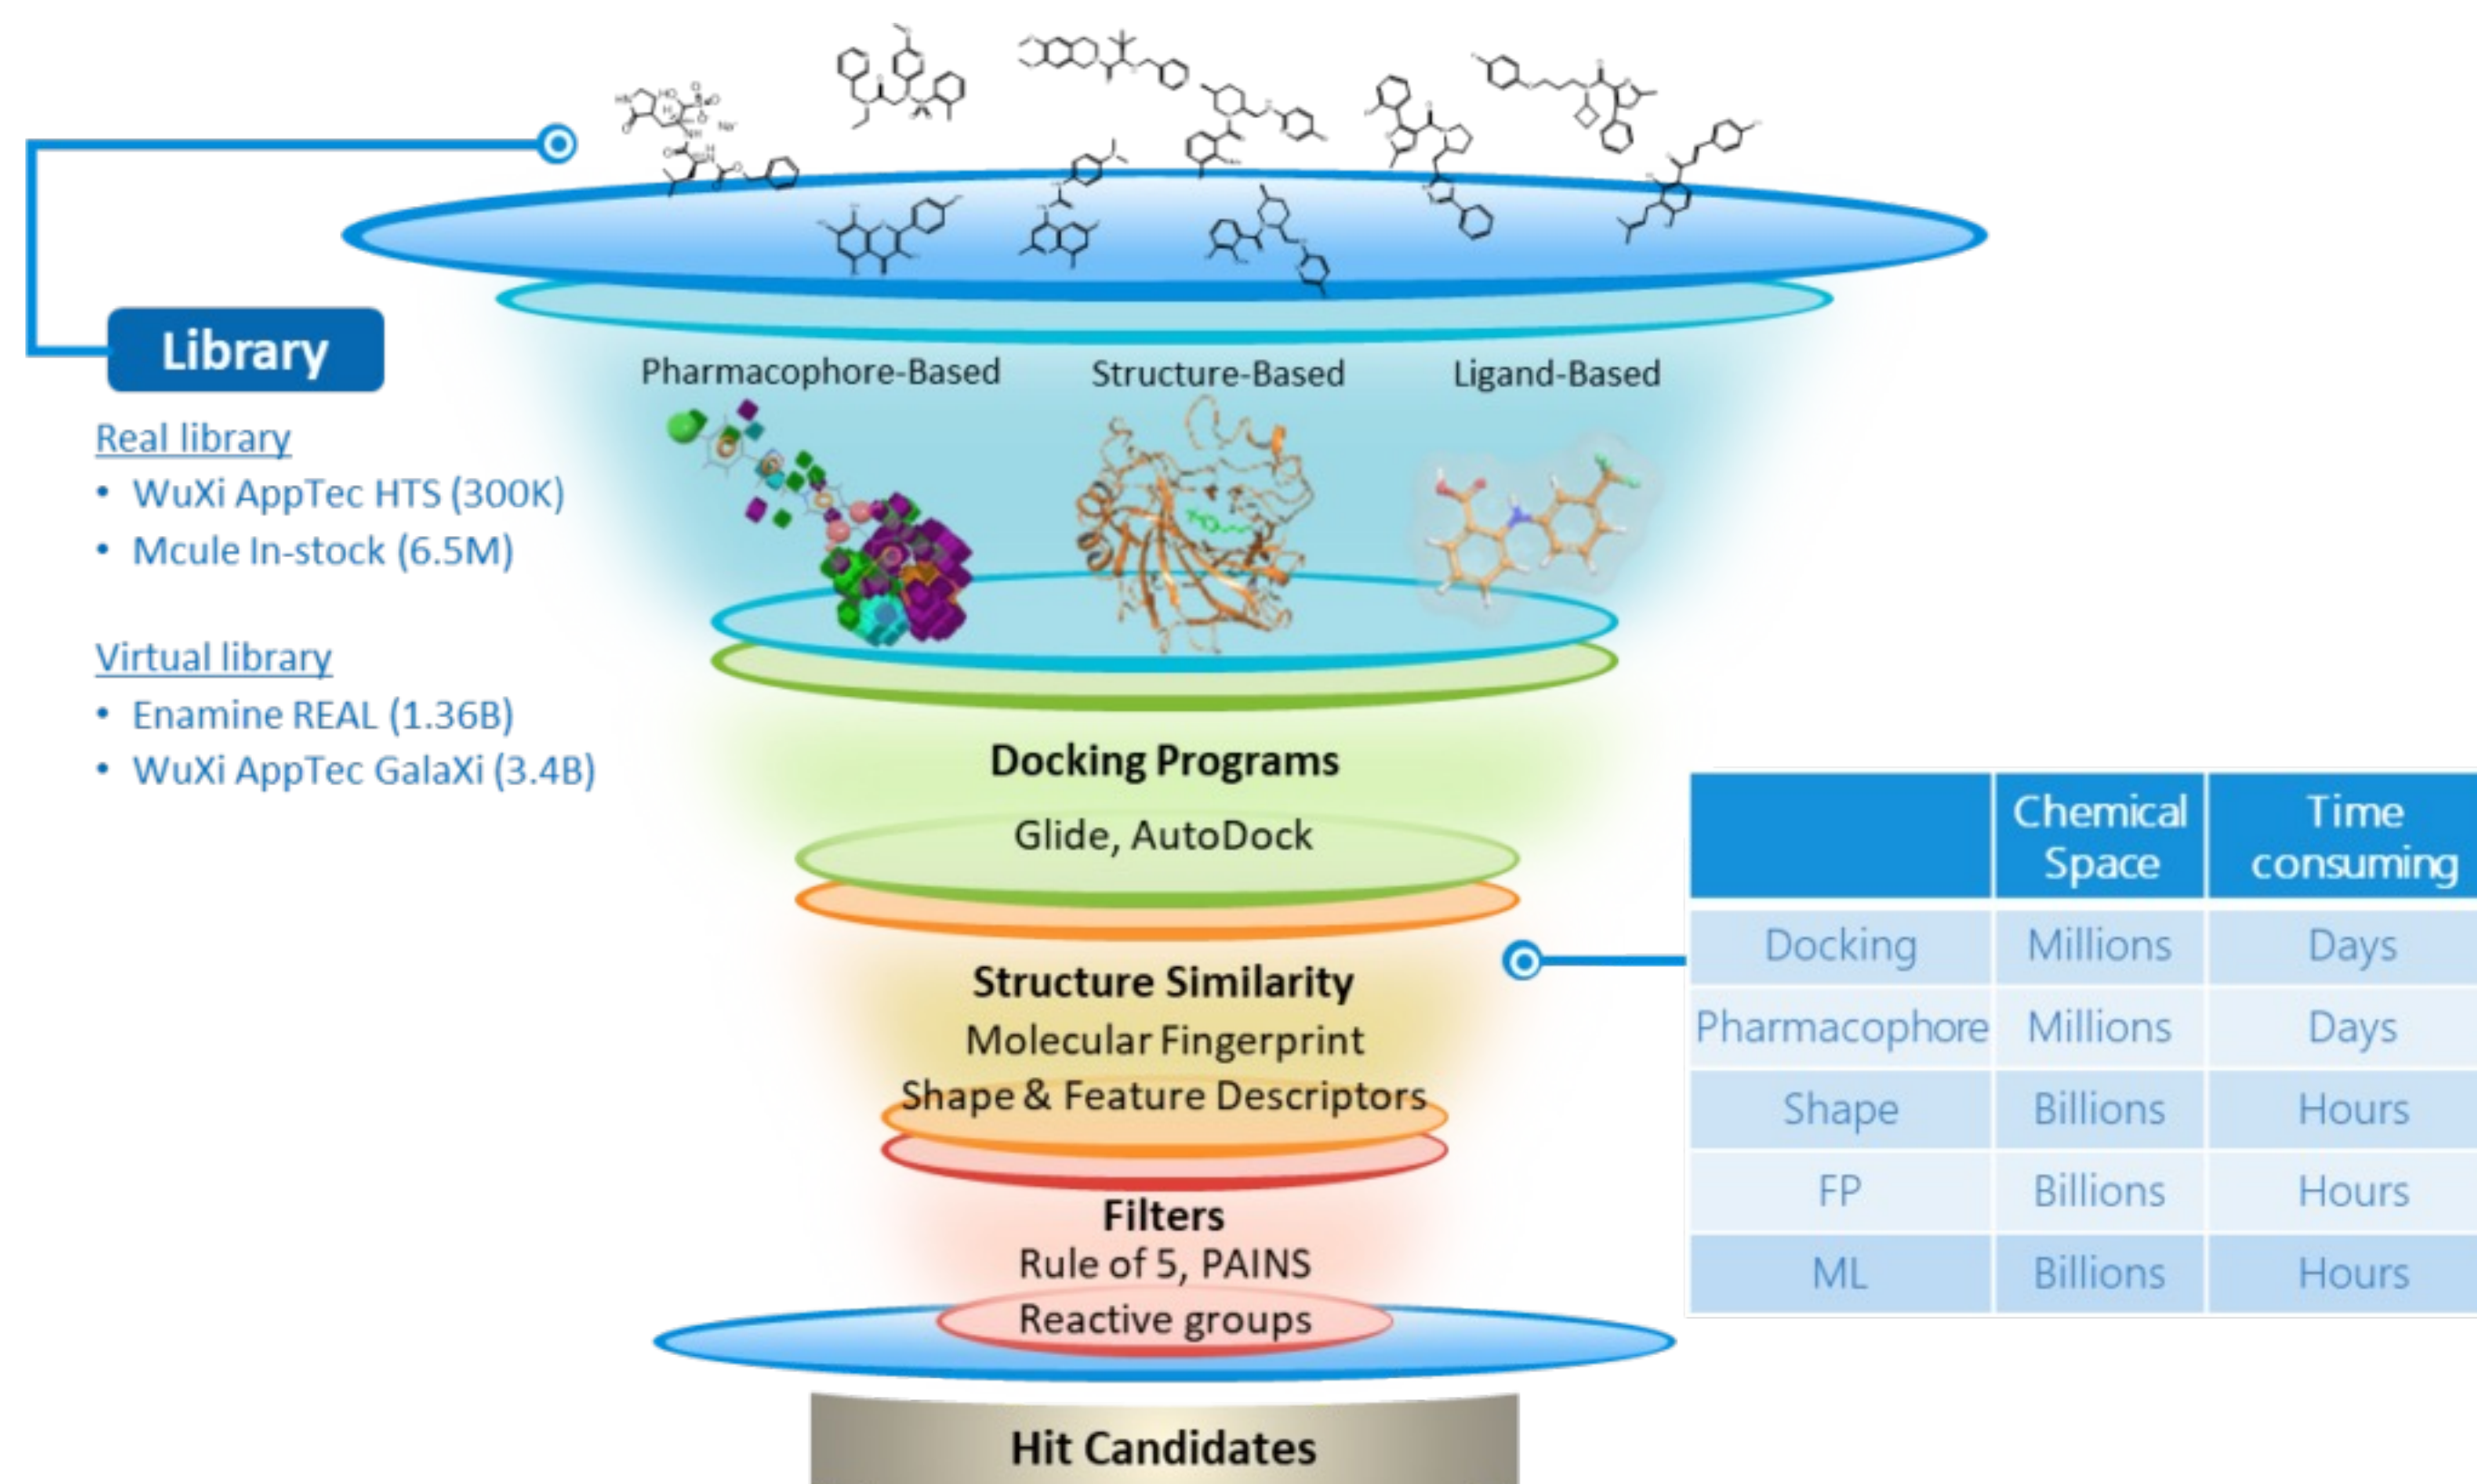
% of assets per pipeline stage overall



1. Drug Design: Current paradigm

Virtual Screening

Efficiently searching large chemical space to find hits



 Central idea:

Search the vast drug-like space using
virtual screening
to identify good starting points

Chemical libraries

New chemical libraries are vast, diverse and commercial available


ZINC

Most common virtual library. ~1.3 billion purchasable molecules by pooling together 310 commercial catalogs




ChEMBL

2.6 million manually curated database of bioactive molecules with drug-like properties



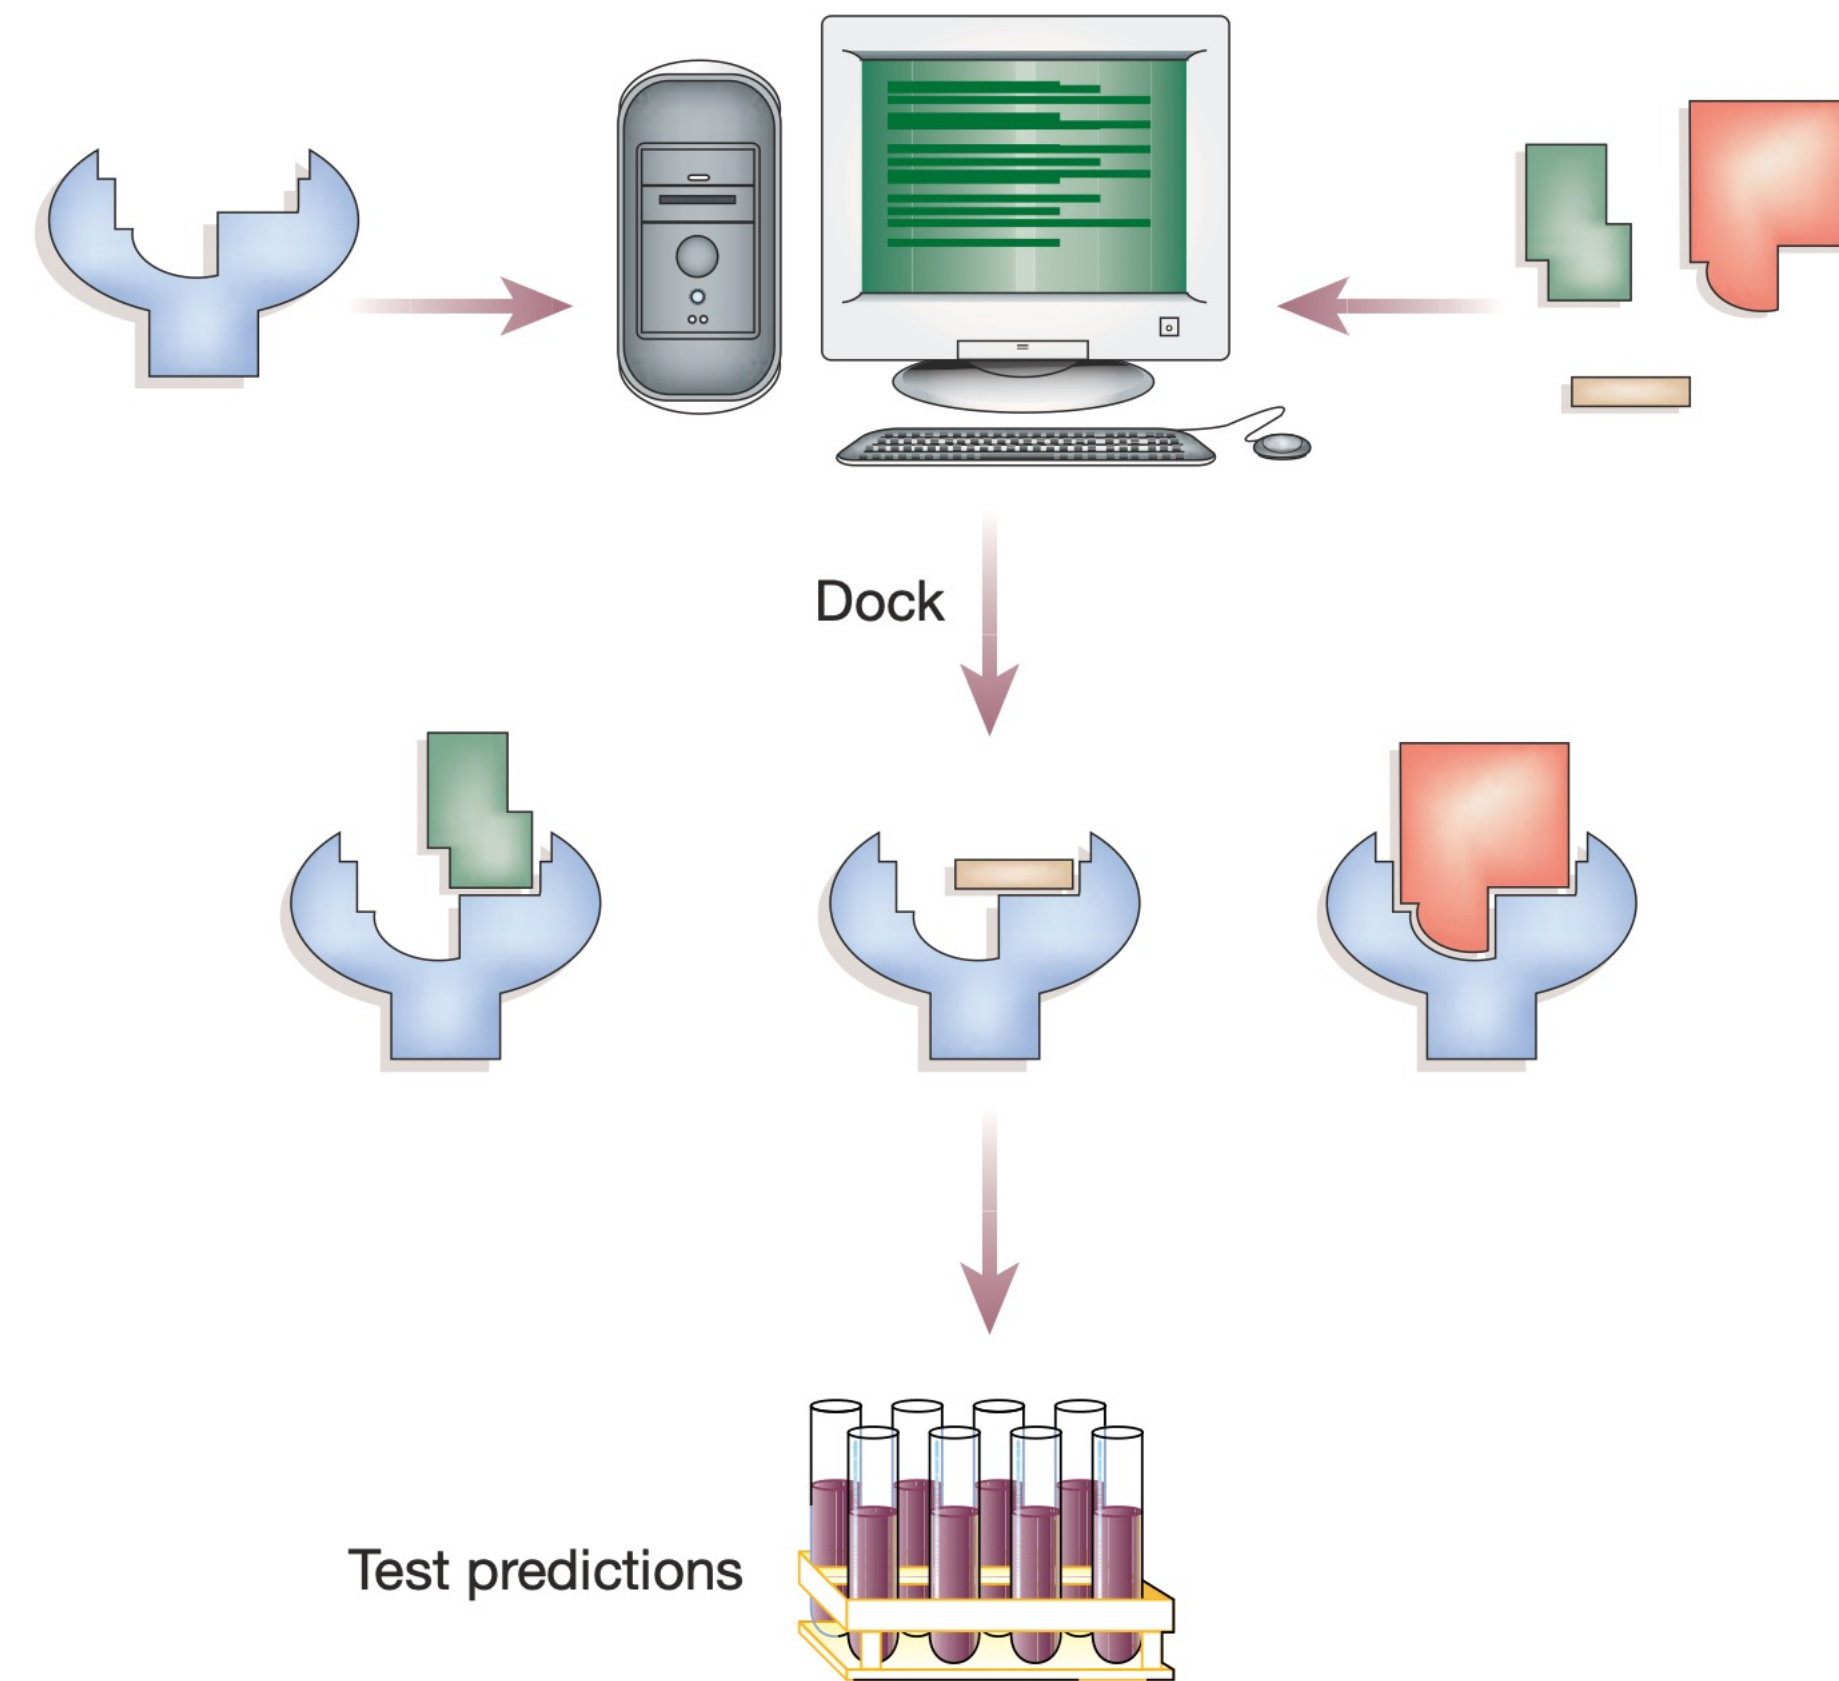
Enamine REAL

(Relatively) small library of 'building-blocks can be combined combinatorically to make a vast chemical space



Protein-ligand docking

Approximates protein-ligand complementarity and affinity



Protein-ligand docking

Combines 2 techniques: a scoring function and optimisation



Optimisation algorithm

Common example is Vina

$$\Delta G_{bind} = \Delta G_{solvent} + \Delta G_{conf} + \Delta G_{int} + \Delta G_{rot} + \Delta G_{t/t} + \Delta G_{vib}$$

Optimisation algorithm

Usually a mix of local and global search



1. DiffDock: Docking with deep learning

Recap: Diffusion Models


Mapping noise back to data



$$L_t^{\text{simple}} = \mathbb{E}_{t \sim [1, T], \mathbf{x}_0, \epsilon_t} \left[\|\epsilon_t - \epsilon_{\theta}(\mathbf{x}_t, t)\|^2 \right]$$

Recap: Diffusion Models


Forward Process = Noising to a reference distribution



$$L_t^{\text{simple}} = \mathbb{E}_{t \sim [1, T], \mathbf{x}_0, \epsilon_t} \left[\|\epsilon_t - \epsilon_{\theta}(\mathbf{x}_t, t)\|^2 \right]$$


Recap: Diffusion Models

Reverse Process: Denoising to our target distribution



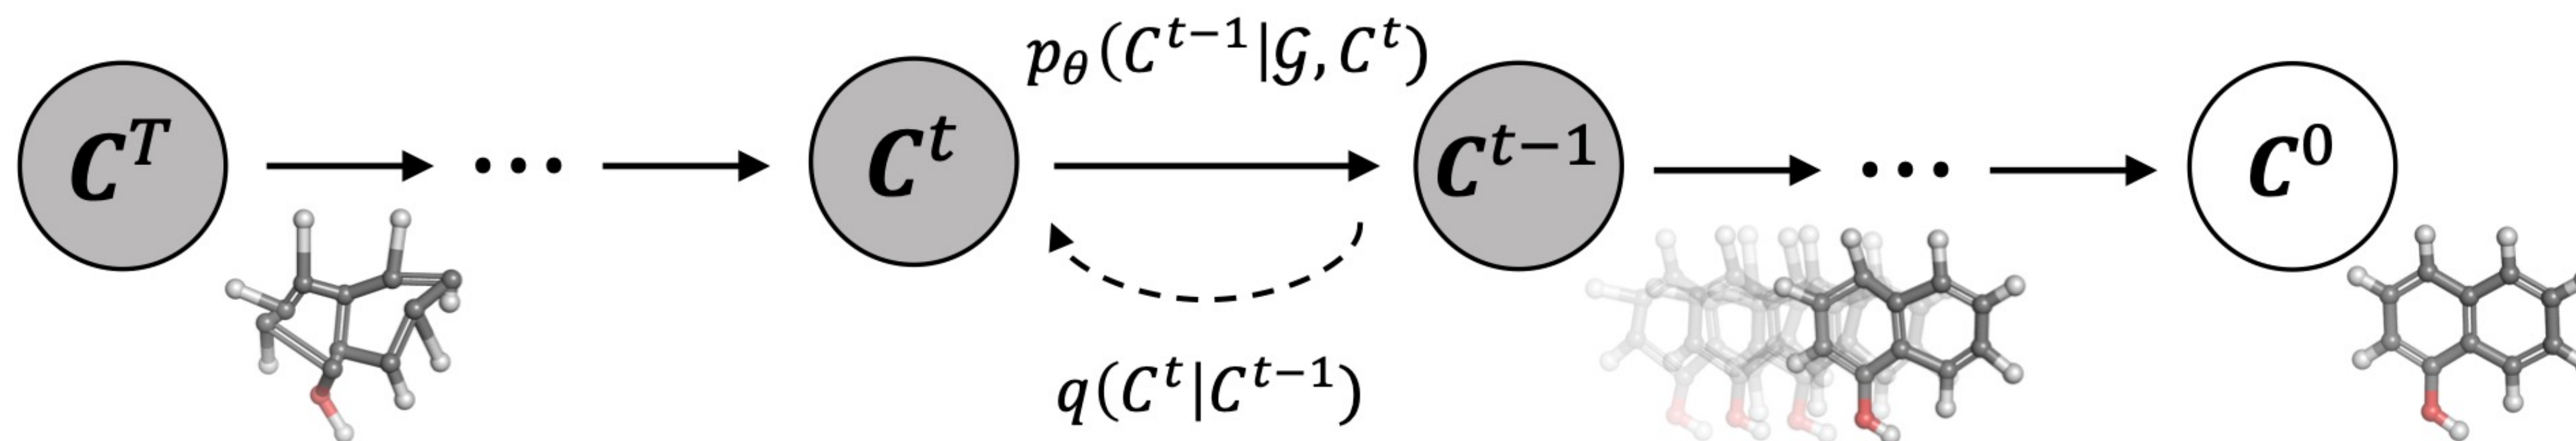
Recap: Geometric Deep Learning

GDL is the application of deep learning to objects that inhabit geometric domains (spaces)




Background

GeoDiff – early work on conformer generation using diffusion models




Torsion Diffusion

Simplified to only generate the degrees of freedom in a molecule (torsion angles)



DiffDock

Docking = Torsional Diffusion + SE(3) diffusion (global rotations and translations)



Method	Holo crystal proteins			
	Top-1 RMSD %<2	Med.	Top-5 RMSD %<2	Med.
GNINA	22.9	7.7	32.9	4.5
SMINA	18.7	7.1	29.3	4.6
GLIDE	21.8	9.3	-	-
EQUIBIND	5.5	6.2	-	-
TANKBIND	20.4	4.0	24.5	3.4
P2RANK+SMINA	20.4	6.9	33.2	4.4
P2RANK+GNINA	28.8	5.5	38.3	3.4
EQUIBIND+SMINA	23.2	6.5	38.6	3.4
EQUIBIND+GNINA	28.8	4.9	39.1	3.1
DIFFDOCK (10)	35.0	3.6	40.7	2.65
DIFFDOCK (40)	38.2	3.3	44.7	2.40

DiffDock

Training by learning translational, rotational and torsional diffusion kernels

Algorithm 1: Training procedure (single epoch)

Input: Training pairs $\{(\mathbf{x}^*, \mathbf{y})\}$, RDKit predictions $\{\mathbf{c}\}$

foreach $\mathbf{c}, \mathbf{x}^*, \mathbf{y}$ **do**

Let $\mathbf{x}_0 \leftarrow \arg \min_{\mathbf{x}^\dagger \in \mathcal{M}_{\mathbf{c}}} \text{RMSD}(\mathbf{x}^*, \mathbf{x}^\dagger)$;

Compute $(\mathbf{r}_0, R_0, \boldsymbol{\theta}_0) \leftarrow A_{\mathbf{c}}^{-1}(\mathbf{x}_0)$;

Sample $t \sim \text{Uni}([0, 1])$;

Sample $\Delta \mathbf{r}, \Delta R, \Delta \boldsymbol{\theta}$ from diffusion kernels $p_t^{\text{tr}}(\cdot | 0), p_t^{\text{rot}}(\cdot | 0), p_t^{\text{tor}}(\cdot | 0)$;

Set $\mathbf{r}_t \leftarrow \mathbf{r}_0 + \Delta \mathbf{r}$;

Set $R_t \leftarrow (\Delta R)R_0$;

Set $\boldsymbol{\theta}_t \leftarrow \boldsymbol{\theta}_0 + \Delta \boldsymbol{\theta} \pmod{2\pi}$;

Compute $\mathbf{x}_t \leftarrow A((\mathbf{r}_t, R_t, \boldsymbol{\theta}_t), \mathbf{c})$;

Predict scores $\alpha \in \mathbb{R}^3, \beta \in \mathbb{R}^3, \gamma \in \mathbb{R}^m = \mathbf{s}(\mathbf{x}_t, \mathbf{c}, \mathbf{y}, t)$;

Take optimization step on loss

$$\mathcal{L} = \|\alpha - \nabla \log p_t^{\text{tr}}(\Delta \mathbf{r} | 0)\|^2 + \|\beta - \nabla \log p_t^{\text{rot}}(\Delta R | 0)\|^2 + \|\gamma - \nabla \log p_t^{\text{tor}}(\Delta \boldsymbol{\theta} | 0)\|^2$$

DiffDock

Training by learning translational, rotational and torsional diffusion kernels

Algorithm 2: Inference procedure

Input: RDKit prediction \mathbf{c} , protein structure \mathbf{y} (both centered at origin)

Output: Sampled ligand pose \mathbf{x}_0

Sample $\boldsymbol{\theta}_N \sim \text{Uni}(SO(2)^m)$, $R_N \sim \text{Uni}(SO(3))$, $\mathbf{r}_N \sim \mathcal{N}(0, \sigma_{\text{tor}}^2(T))$;

Let $\mathbf{x}_N = A((\mathbf{r}_N, R_N, \boldsymbol{\theta}_N), \mathbf{c})$;

for $n \leftarrow N$ **to** 1 **do**

Let $t = n/N$ and $\Delta\sigma_{\text{tr}}^2 = \sigma_{\text{tr}}^2(n/N) - \sigma_{\text{tr}}^2((n-1)/N)$ and similarly for $\Delta\sigma_{\text{rot}}^2, \Delta\sigma_{\text{tor}}^2$;

Predict scores $\alpha \in \mathbb{R}^3, \beta \in \mathbb{R}^3, \gamma \in \mathbb{R}^m \leftarrow \mathbf{s}(\mathbf{x}_n, \mathbf{c}, \mathbf{y}, t)$;

Sample $\mathbf{z}_{\text{tr}}, \mathbf{z}_{\text{rot}}, \mathbf{z}_{\text{tor}}$ from $\mathcal{N}(0, \Delta\sigma_{\text{tr}}^2), \mathcal{N}(0, \Delta\sigma_{\text{rot}}^2), \mathcal{N}(0, \Delta\sigma_{\text{tor}}^2)$ respectively;

Set $\mathbf{r}_{n-1} \leftarrow \mathbf{r}_n + \Delta\sigma_{\text{tr}}^2\alpha + \mathbf{z}_{\text{tr}}$;

Set $R_{n-1} \leftarrow \mathbf{R}(\Delta\sigma_{\text{rot}}^2\beta + \mathbf{z}_{\text{rot}})R_n$;


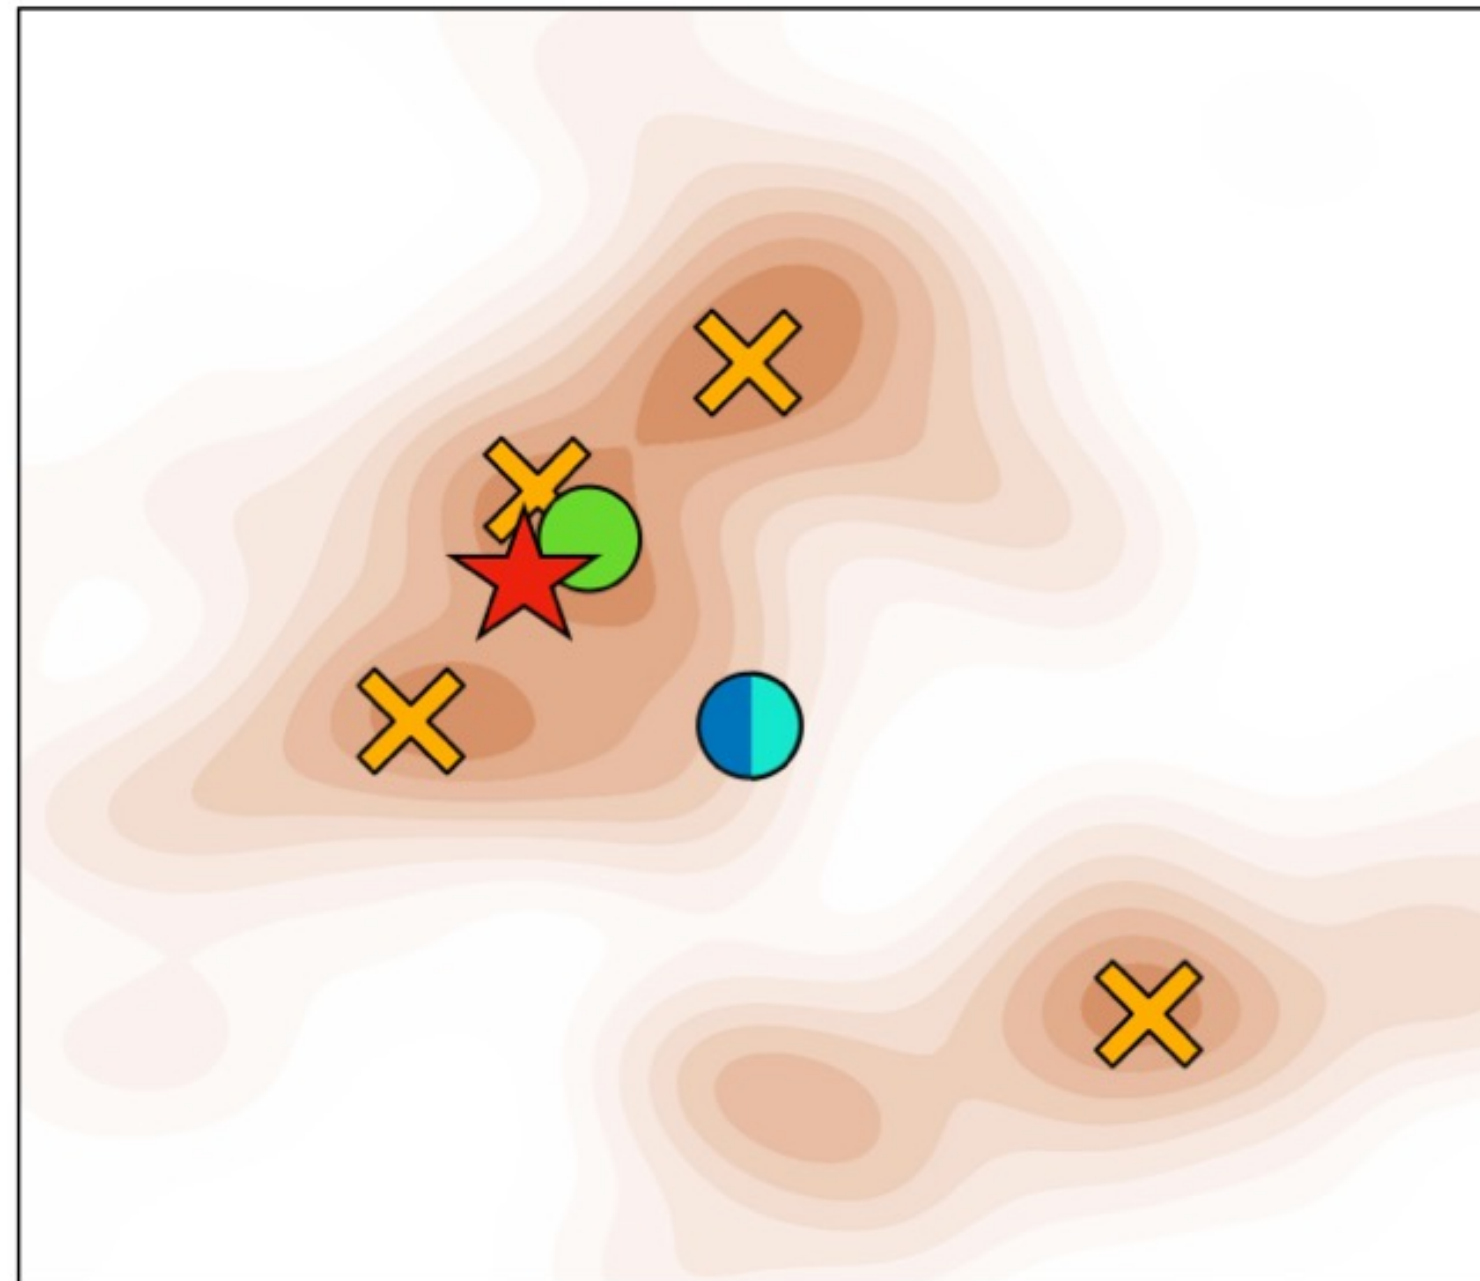
Set $\boldsymbol{\theta}_{n-1} \leftarrow \boldsymbol{\theta}_n + (\Delta\sigma_{\text{tor}}^2\gamma + \mathbf{z}_{\text{tor}}) \bmod 2\pi$;

Compute $\mathbf{x}_{n-1} \leftarrow A((\mathbf{r}_{n-1}, R_{n-1}, \boldsymbol{\theta}_{n-1}), \mathbf{c})$;

Return \mathbf{x}_0 ;

DiffDock

In theory, generative modelling allows us to approximate and sample from the whole binding landscape




Limitations of DL docking are significant

Many substantial issues can be masked when only measuring performance by RMSD

Poor evaluations


DL-based methods are often evaluated using blind docking, something conventional methods are not designed for



- When evaluated by docking molecules into known pockets, the traditional methods still outperform DiffDock
- DiffDock is actually a SOTA binding pocket prediction algorithm

Unrealistic molecules

DL-based methods can have significant biophysical violations, even if docking RMSD is low

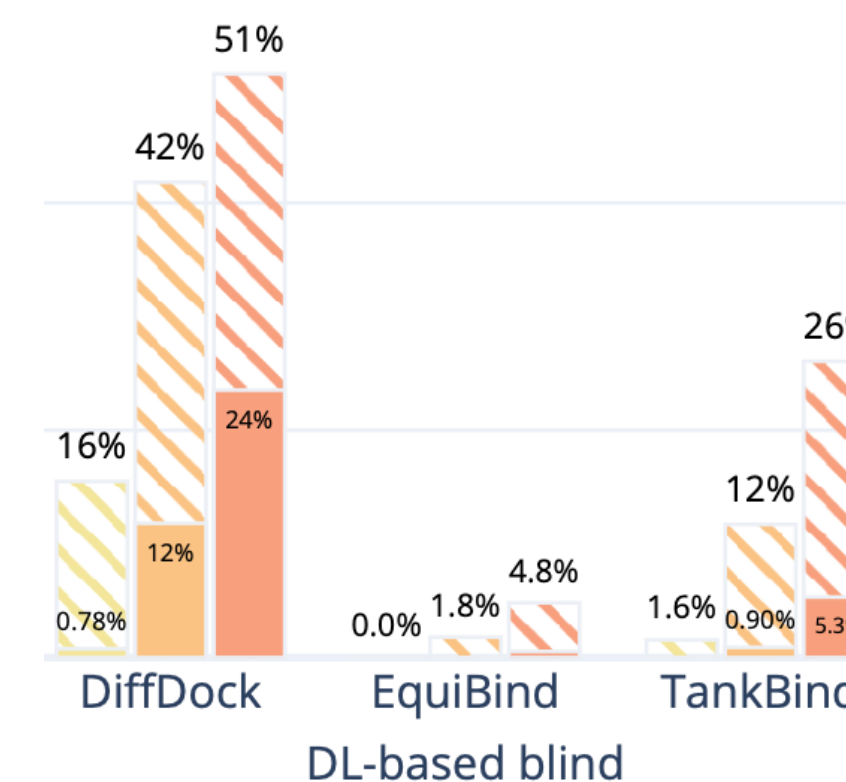


Recommended Reading:

- Martin Butterschoen– [PoseBusters](#)

Poor generalisation

DL-based methods struggle to generalise to proteins not seen during training



- Generalisation to novel receptors and molecules is essential
- Issue also large in molecule scaffold generalisation but less studied


Recommended Reading:

- Martin Butterschoen– [PoseBusters](#)

2. SBDD with Generative Models

SBDD with Generation Models

Rephrasing SBDD as learning a conditional probability distribution



$$\text{SBDD} = p(\text{molecule}|\text{receptor})$$

SBDD with Generation Models

Rephrasing SBDD as learning a conditional probability distribution


$$\text{SBDD} = p(\text{molecule}|\text{receptor})$$

💡 Central idea:

Treat drug design as a
condition generation problem
by learning from data


SBDD with Diffusion Models

Learning to generate complimentary molecules in 3D




Generative Modelling for Molecule Generation

1. All-at-once (one-shot)



(Zang et al, 2020)




Node independence
assumed



Arbitrary node ordering

2. Node-by-node (autoregressive)




(Imrie et al, 2020)




Different generation
traces not equal

Generative Modelling for Molecule Generation

1. All-at-once (one-shot)



(Zang et al, 2020)




Node independence
assumed



Arbitrary node ordering

2. Node-by-node (autoregressive)



(Imrie et al, 2020)




Different generation
traces not equal


Generative Modelling for Molecule Generation



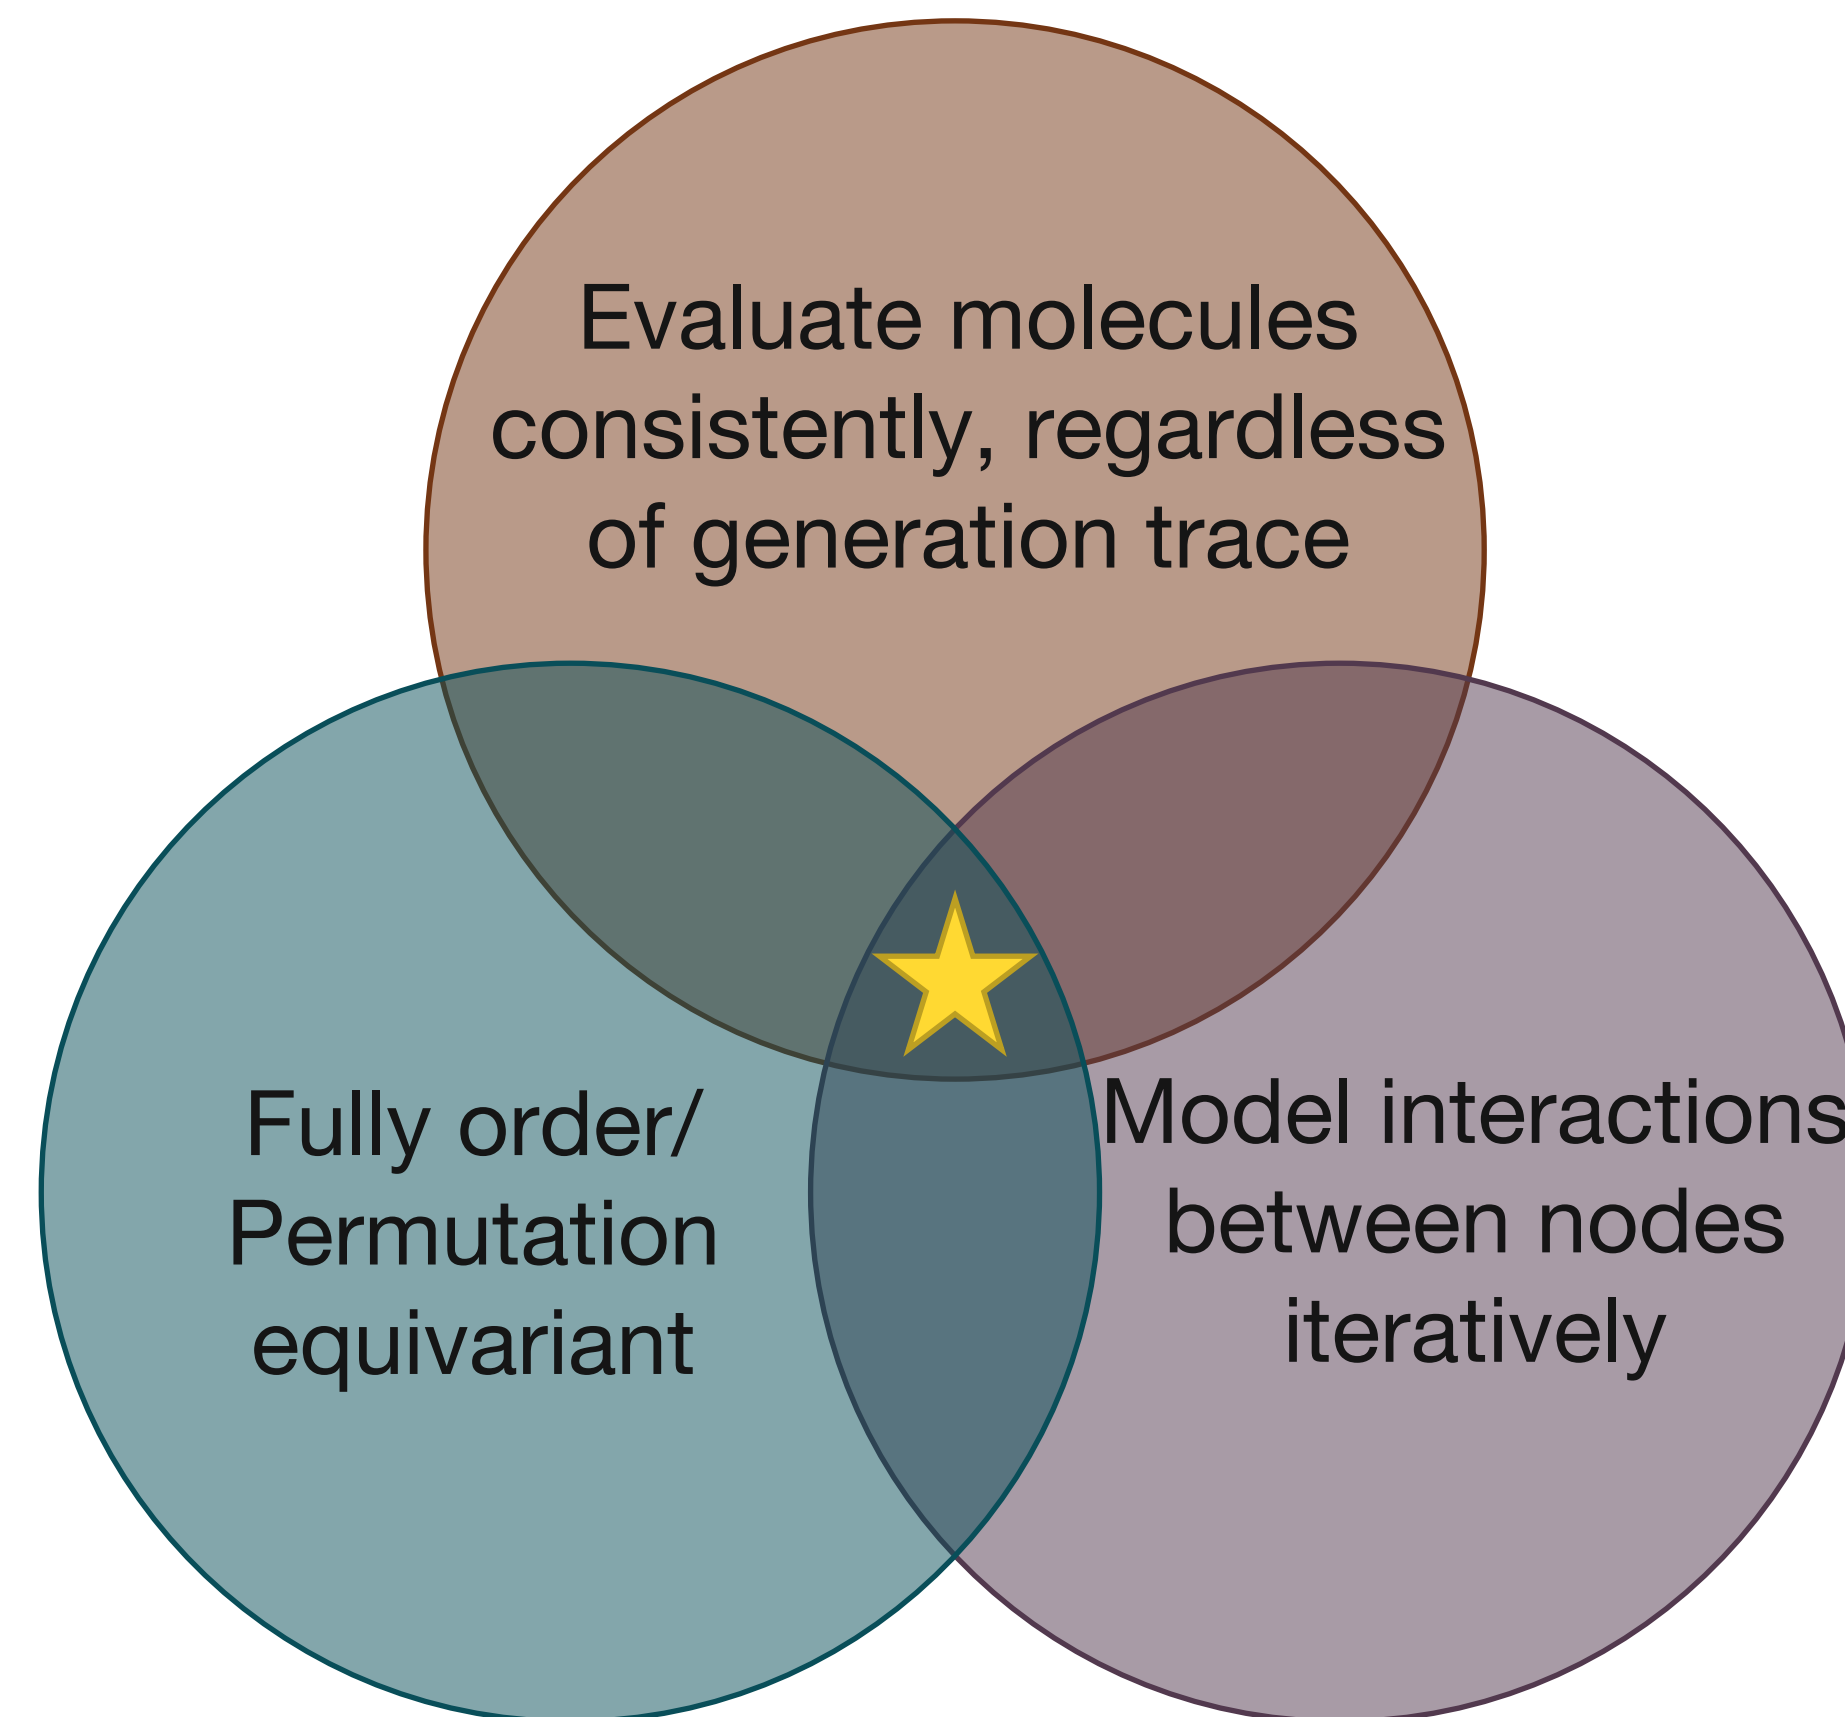
Node independence
assumed



Arbitrary node ordering



Different generation
traces not equal



Structure-based Drug Design with Equivariant Diffusion Models

Arne S. (EPFL), Yuanqi Du (Cornell), Charles Harris (Cambridge), Arian J. (Cambridge), Ilya Igashov (EPFL), Weitao Du (Cornell), Tom Blundell (Cambridge), Pietro Lio (Cambridge), Carla Gomes (Cornell), Max Welling (Amsterdam/Microsoft), Michael Bronstein (Oxford/Twitter), Bruno Correia (EPFL)






 UNIVERSITY OF
CAMBRIDGE

 EPFL

 Cornell University

 UNIVERSITY OF
OXFORD

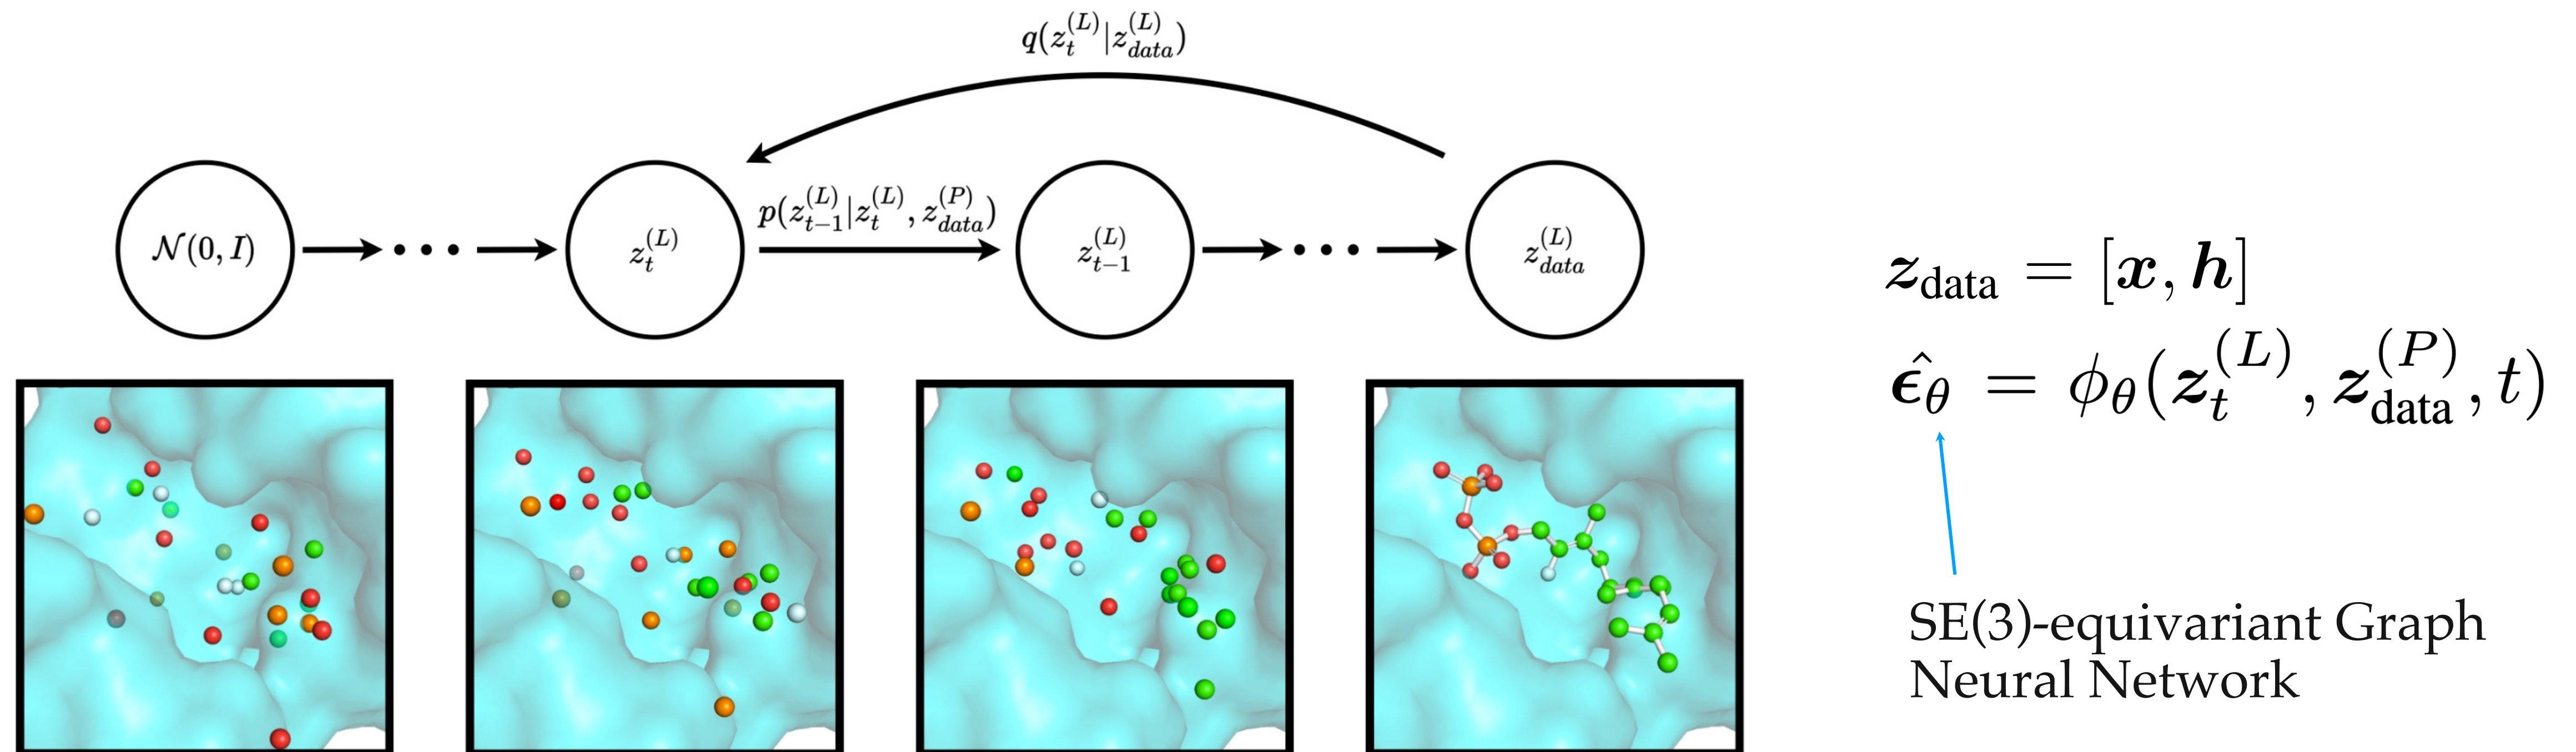
 Microsoft Research



DiffSBDD:



A Diffusion Model for Structure-based Drug Design

- Both proteins and ligands are represented as all-atom graphs
- Learns the transitional probability distribution $p_{\theta} \left(\mathbf{z}_{t-1}^{(L)} \mid \mathbf{z}_t^{(L)}, \mathbf{z}_{data}^{(P)} \right)$
- Denoising network $\hat{\epsilon}_{\theta}$ constructs samples





DiffSBDD: Results

Conditional (2jjg)





Vina: -6.5 Sim: 0.27 Vina: -6.7 Sim: 0.24 Vina: -6.6 Sim: 0.21
 QED: 0.49 SA: 0.43 QED: 0.63 SA: 0.35 QED: 0.54 SA: 0.27

Inpainting-Ca (2jjg)





Vina: -6.5 Sim: 0.27 Vina: -6.3 Sim: 0.19 Vina: -6.4 Sim: 0.19
 QED: 0.44 SA: 0.29 QED: 0.53 SA: 0.35 QED: 0.21 SA: 0.35

Reference (2jjg)





Vina: -5.9 Sim: 1
 QED: 0.56 SA: 0.78

Conditional (3kc1)





Vina: -8.1 Sim: 0.44 Vina: -7.2 Sim: 0.50 Vina: -8.5 Sim: 0.40
 QED: 0.70 SA: 0.45 QED: 0.65 SA: 0.45 QED: 0.63 SA: 0.35

Inpainting-Ca (3kc1)



Vina: -6.9 Sim: 0.40 Vina: -6.9 Sim: 0.32 Vina: -6.4 Sim: 0.23
 QED: 0.15 SA: 0.36 QED: 0.67 SA: 0.27 QED: 0.45 SA: 0.40



Reference (3kc1)



Vina: -6.5 Sim: 1
 QED: 0.72 SA: 0.66



DiffSBDD: Results

Conditional-Ca (6c0b)





Vina: -12.8 Sim: 0.05 Vina: -11.9 Sim: 0.12 Vina: -11.5 Sim: 0.06
 QED: 0.74 SA: 0.45 QED: 0.66 SA: 0.25 QED: 0.68 SA: 0.25

Inpainting-Ca (6c0b)



Vina: -12.4 Sim: 0.07 Vina: -12.3 Sim: 0.07 Vina: -12.2 Sim: 0.12
 QED: 0.76 SA: 0.24 QED: 0.85 SA: 0.25 QED: 0.63 SA: 0.34

Reference (6c0b)




Vina: -8.40 Sim: 1
 QED: 0.36 SA: 0.89

Other Strategies in SBDD

Not having to design molecule de novo simplifies the process

Fragment-based Drug Design


Identifying small molecular motifs (**fragments**) that bind inside a pocket and then **linking** these for form a whole lead molecule



Example: **DiffLinker**

Identifying small molecular motifs (**fragments**) that bind inside a pocket and then **linking** these for form a whole lead molecule

Scaffold Hopping

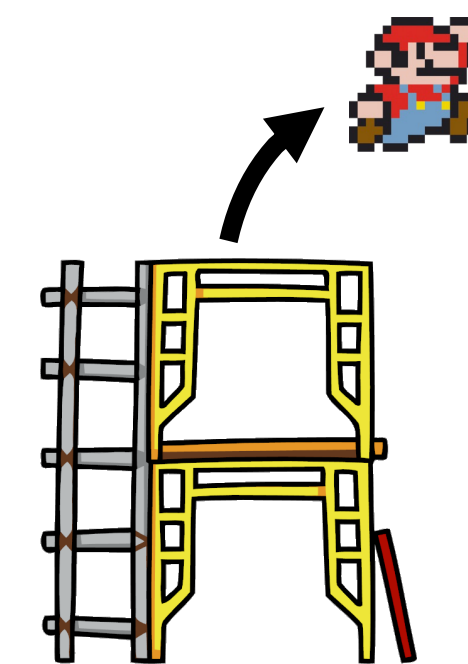



Example: **DiffHopp**

Identifying small molecular motifs (**fragments**) that bind inside a pocket and then **linking** these for form a whole lead molecule


What is scaffold hopping?

Similar properties, novel topology



DiffHopp


A conditional diffusion model for scaffold hopping



Scaffold diversity analysis

Generated scaffold are as diverse as the PDB


- **Can generate diverse scaffolds, regardless of the starting chemotype**
- **However, scaffold space is very large and we are limited by the PDB (<40,000)**



Scaffold diversity analysis

Generated scaffold are as diverse as the PDB

- **Can generate diverse scaffolds, regardless of the starting chemotype**
- **However, scaffold space is very large and we are limited by the PDB (<40,000)**



There are many sub-tasks within SBDD

e.g. Fragment-linking, scaffold hopping




Conditional models are trained on synthetic dataset



Specialist models cannot generalise to new tasks




Need to prespecify atom attachment points




(Imrie et al, 2020)


Limitation of conditional models



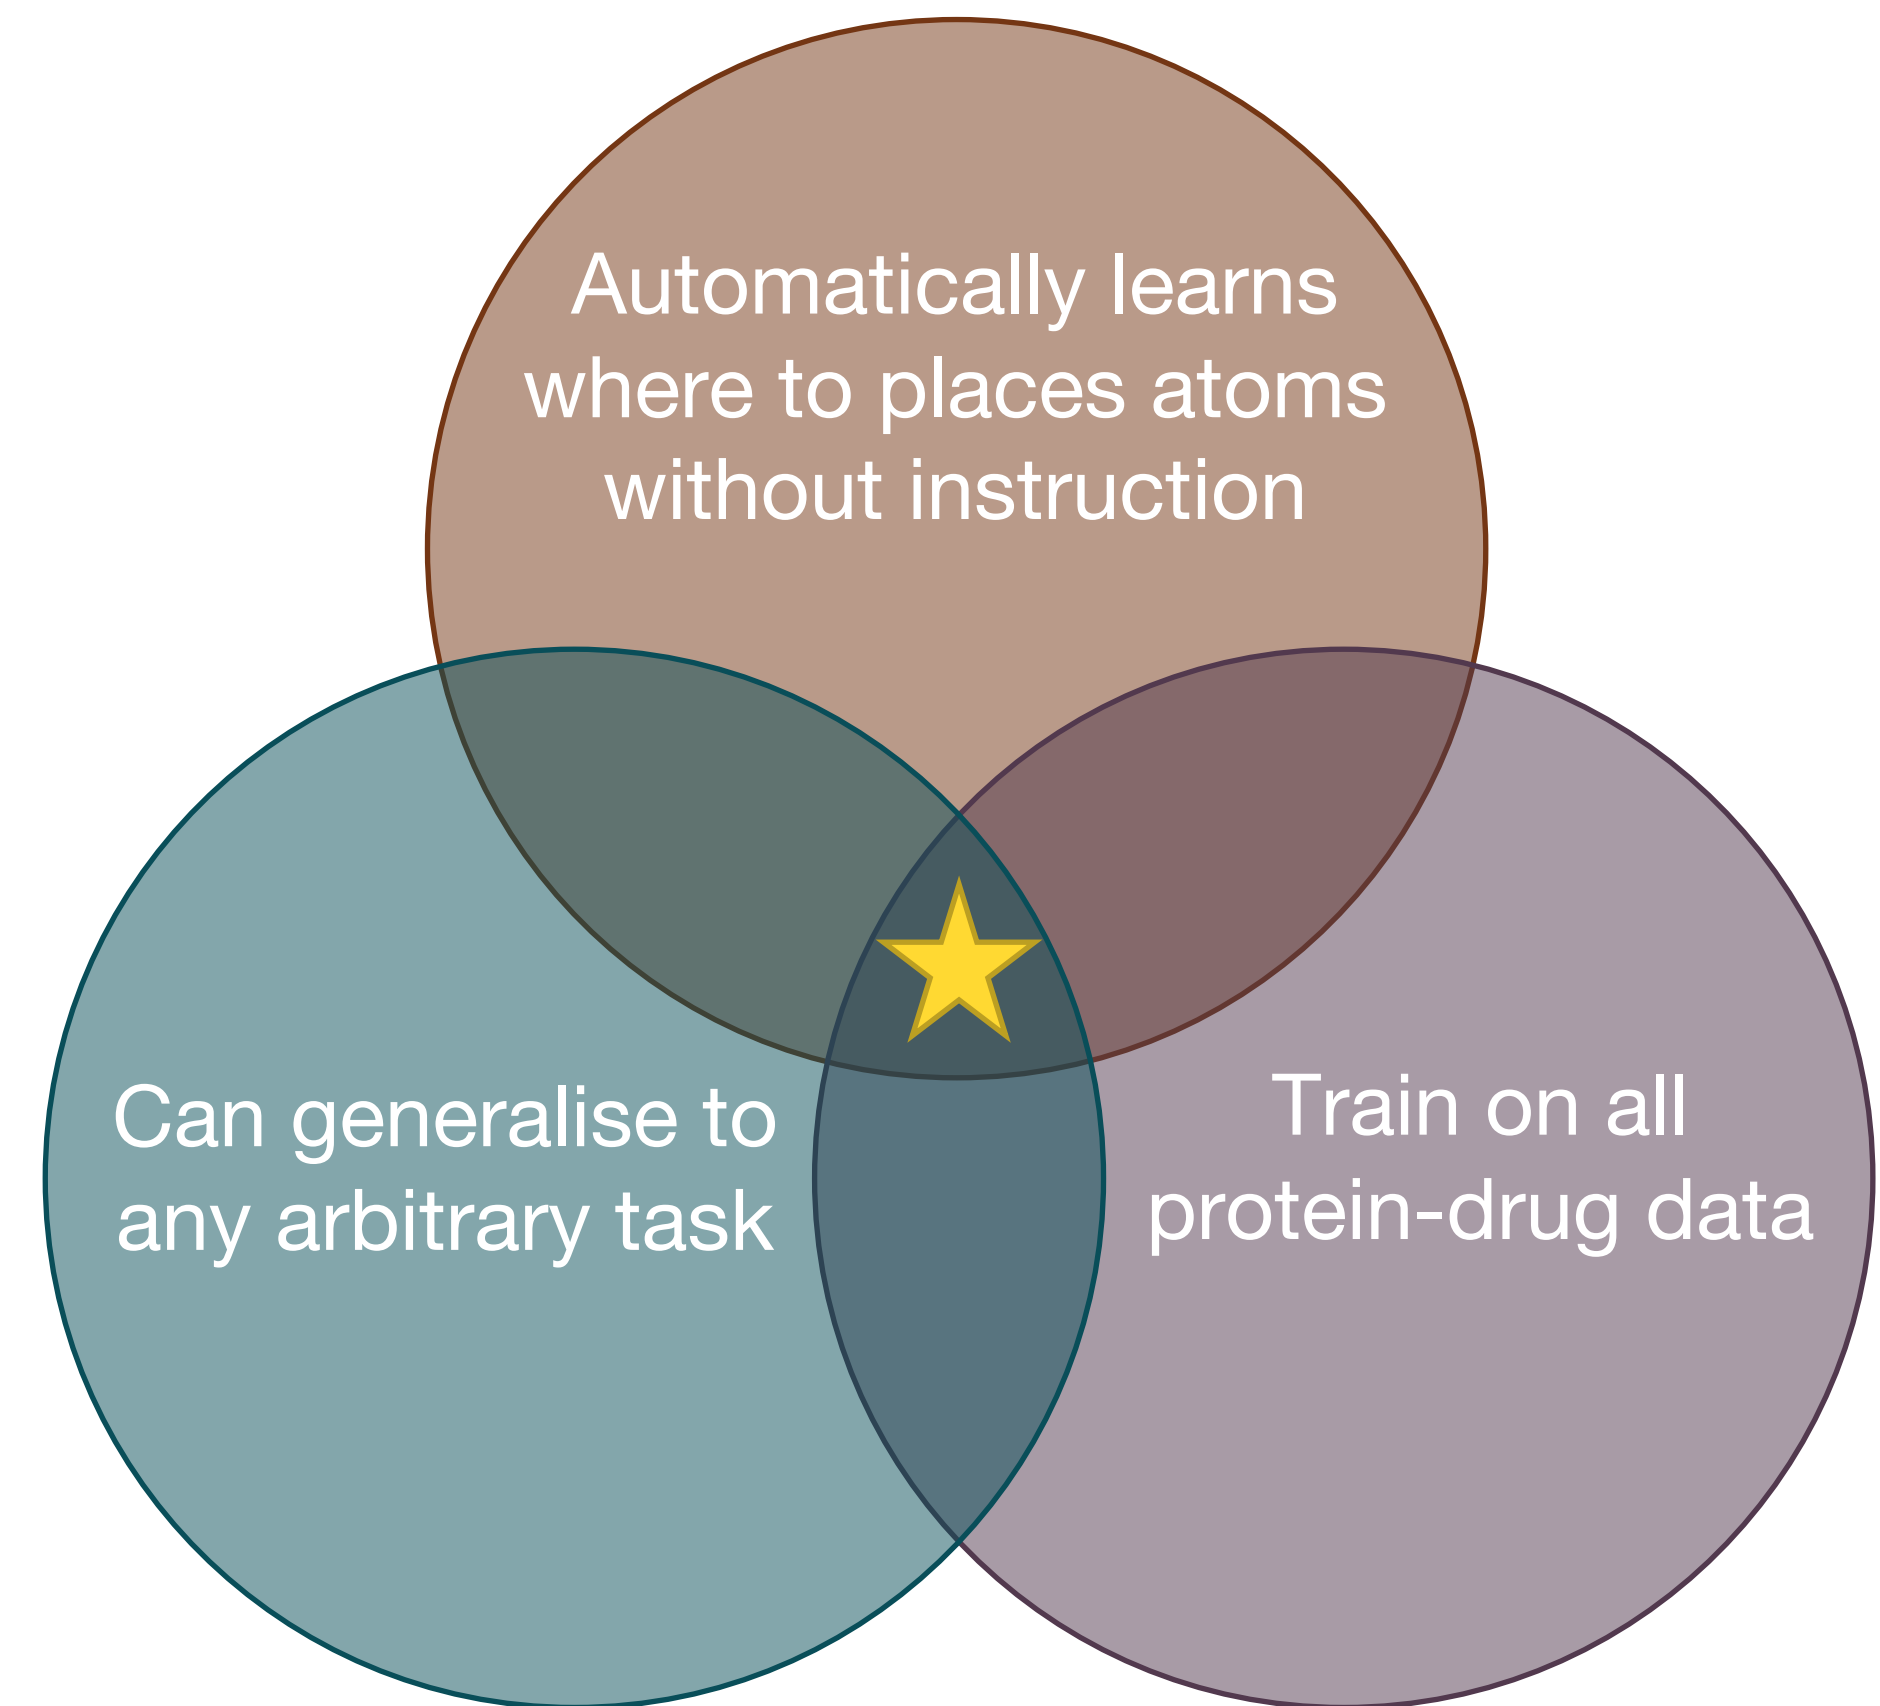
Conditional models are trained on synthetic dataset



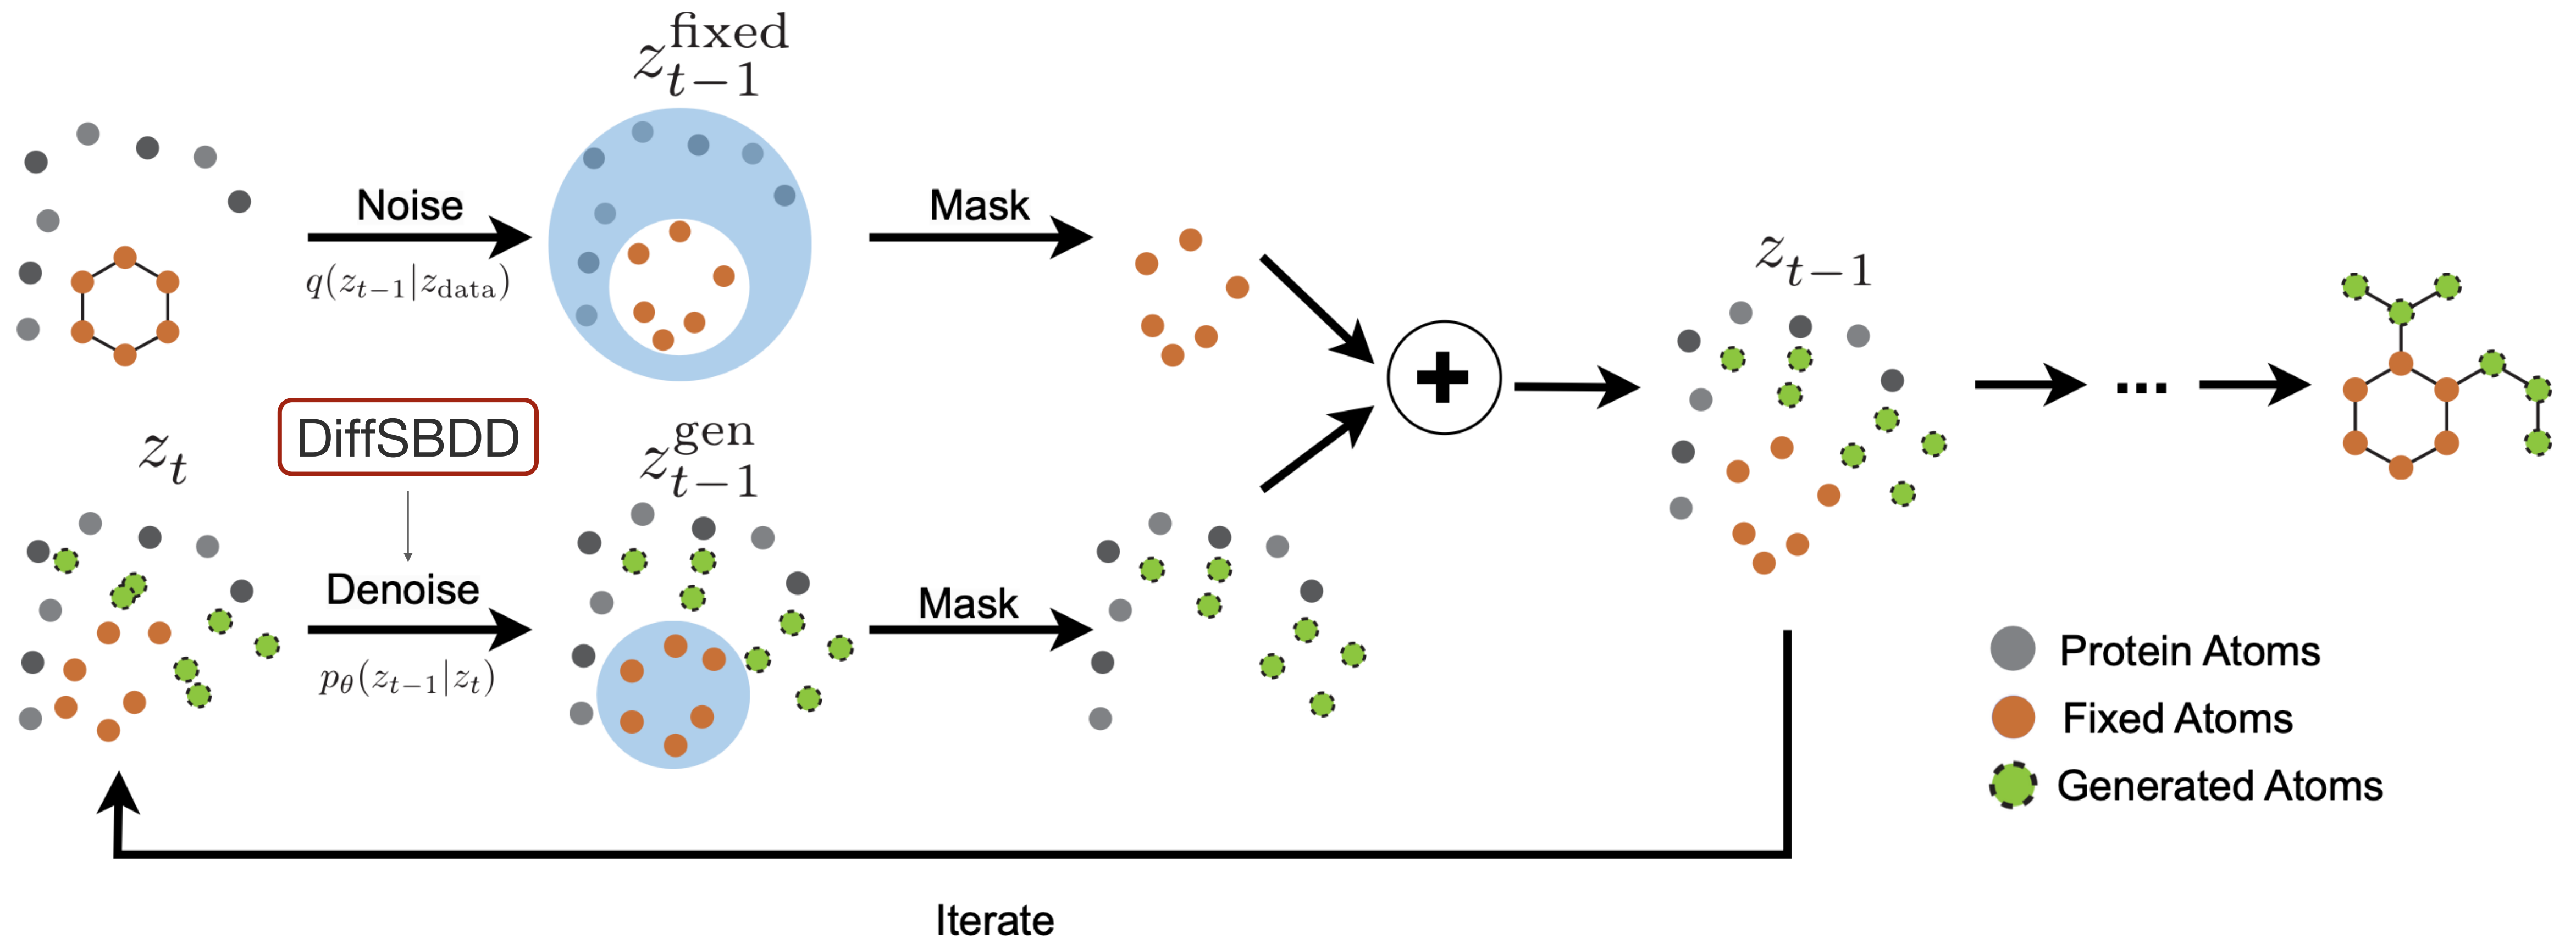
Specialist models cannot generalise to new tasks



Need to prespecify atom attachment points



Molecular inpainting with DiffSBDD





Takeaway



Generative modelling **holds promise** for designing novel drugs but has **no real-world validation**