ETH zürich ETH AI CENTER

Caltech

Teaching Scientific Generators to Reach Discovery Goals

Riccardo De Santi

PhD Candidate at ETH Zurich

@r_de_santi rdesanti@ethz.ch

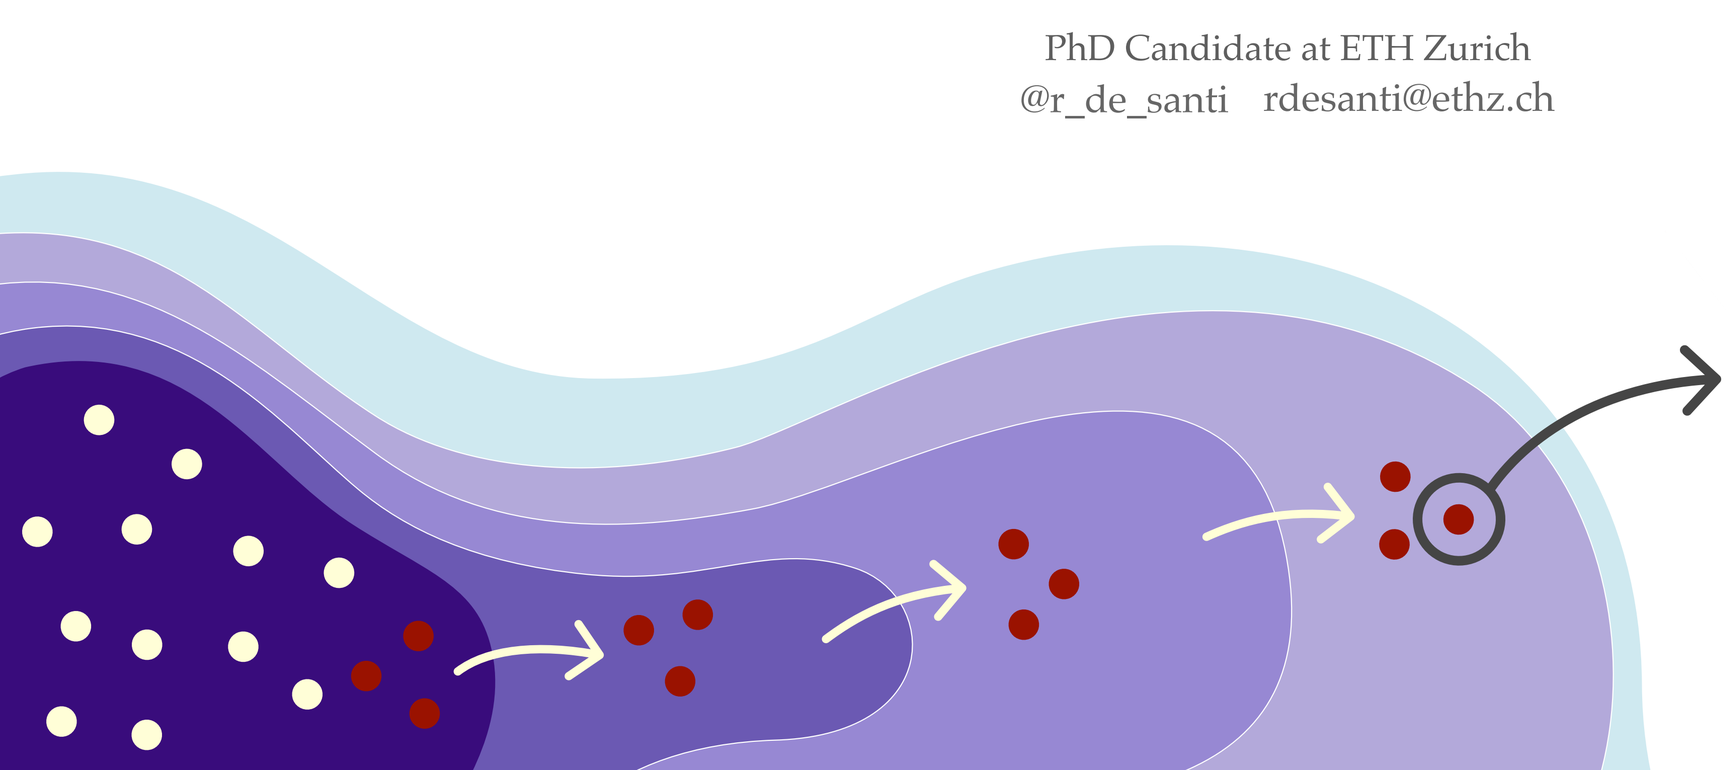

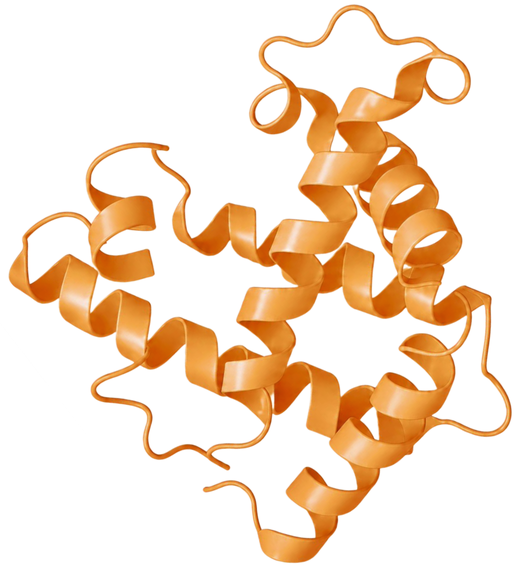

AstraZeneca, November 23rd, 2026

Team



Riccardo De Santi



Ya-Ping Hsieh



Kimon Protopapas



Bruce Lee



Cristian P. Jensen



Zebang Shen



Marin Vlastelica



Zifan Wang



Sophia Tong

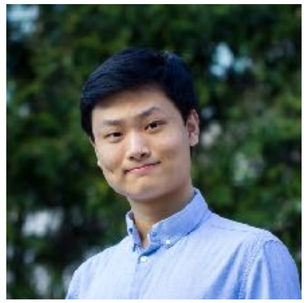

Chenghao Liu



Federico Di Gennaro



Angelo Gnazzo

Advisors



Pranam Chatterjee



Niao He



İsmail İlkan Ceylan



Michael M. Zavlanos

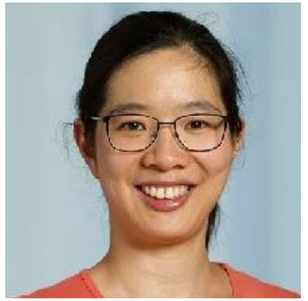

Fanny Yang



Karl H. Johansson



Frances H. Arnold



Yisong Yue



Michael M. Bronstein



Andreas Krause

What is a Discovery Problem?

A real-world discovery task:

> **GOAL**: *"Discover novel **peptides** with promising **anticancer** activity"*

What is a Discovery Problem?

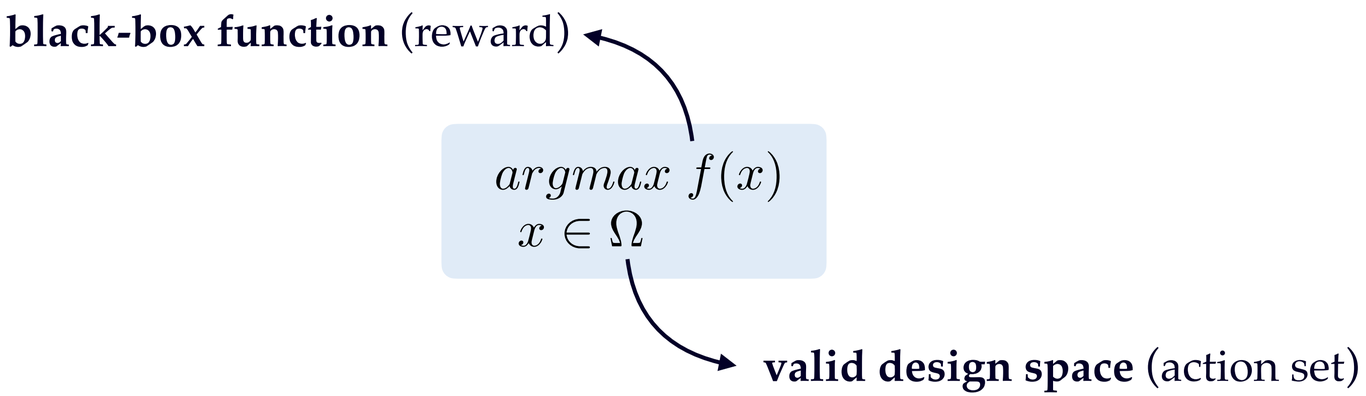

What is a Discovery Problem?

black-box function (reward)

$$\underset{x \in \Omega}{argmax}\ f(x)$$

valid design space (action set)

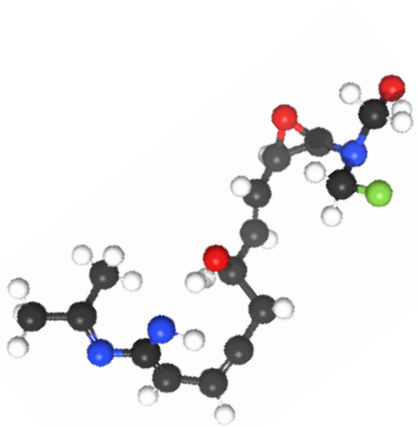
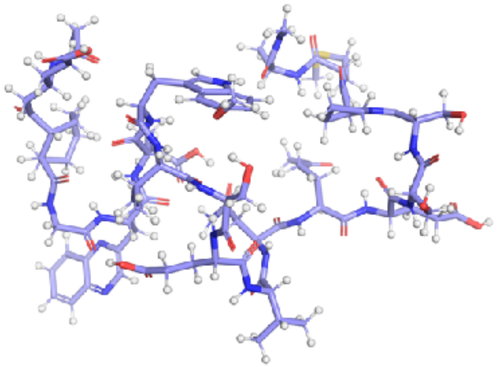
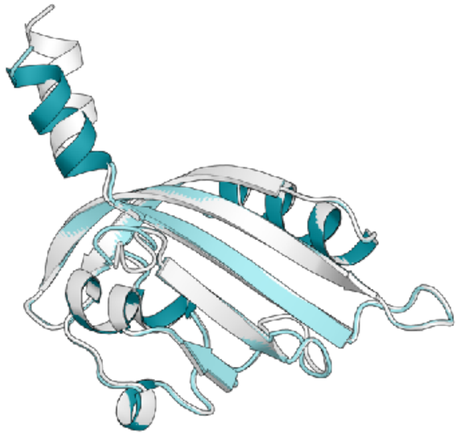
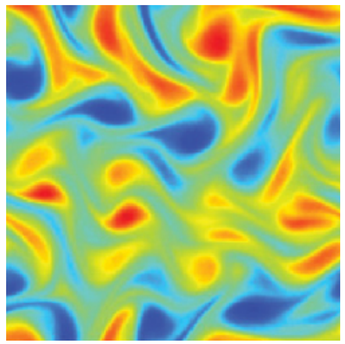
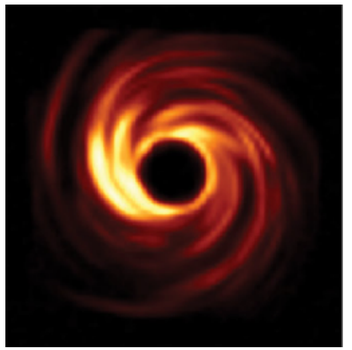

[Zheng et al., 2025]

What is a Discovery Problem?

black-box function (reward)

$$argmax\ f(x)$$
$$x \in \Omega$$

valid design space (action set)

$\approx 10^{60}$ drug discovery

$\approx 10^{80}$ material design

f and Ω are **unknown** and **hard to express**

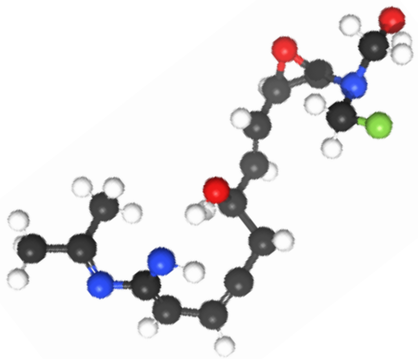

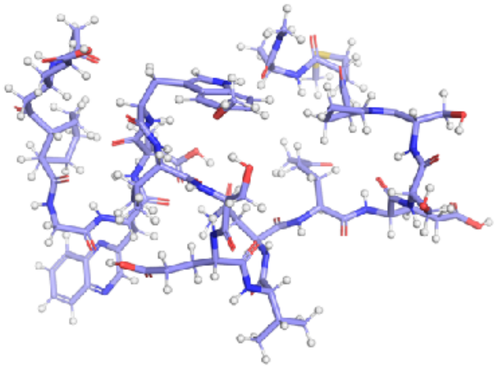

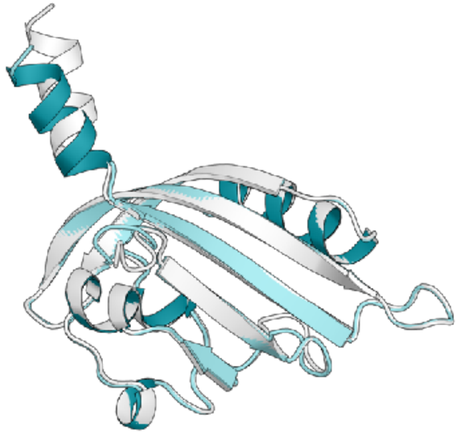

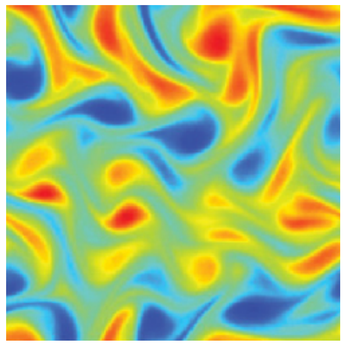

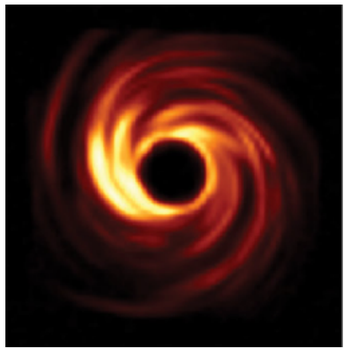

[Zheng et al., 2025]

What is a Discovery Problem?

black-box function (reward)

$$argmax\ f(x)$$
$$x \in \Omega$$

valid design space (action set)

$\approx 10^{60}$ drug discovery

$\approx 10^{80}$ material design

f and Ω are **unknown** and **hard to express**

→ Generative Discovery

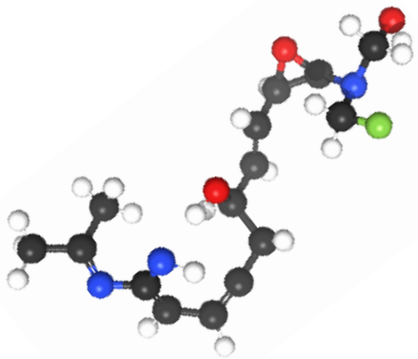

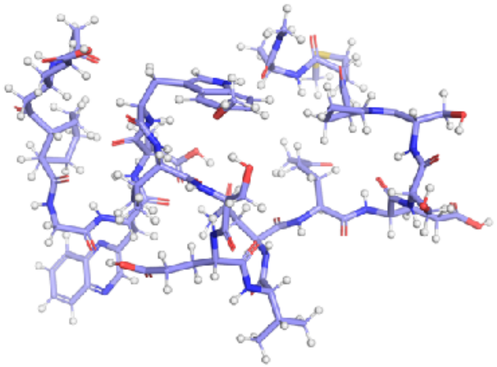

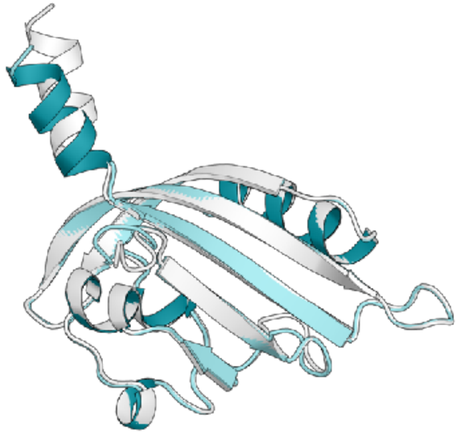

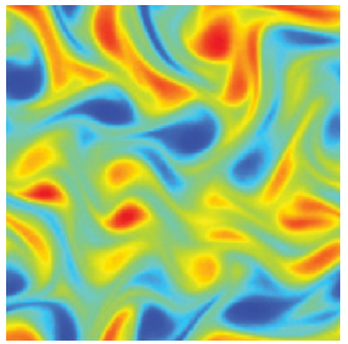

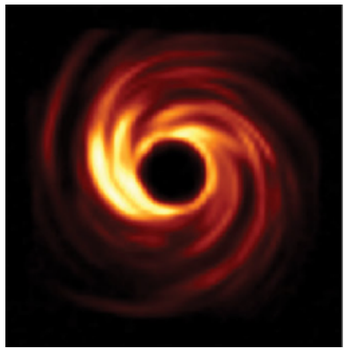

[Zheng et al., 2025]

Generative Discovery Paradigm

Data $D = \{x_i\}_{i=1}^{n}$

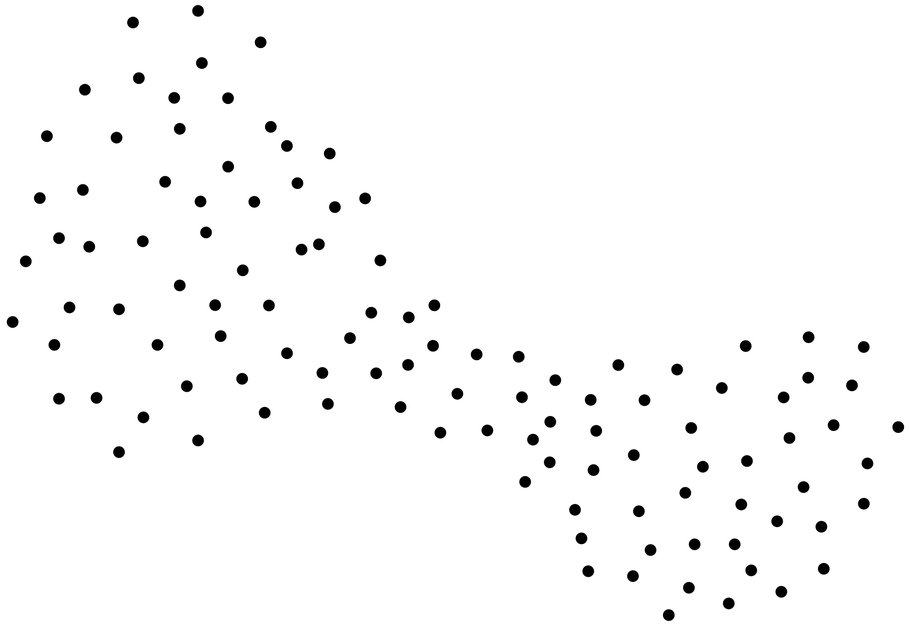

Generative Discovery Paradigm

valid design space Ω (e.g., protein structures)

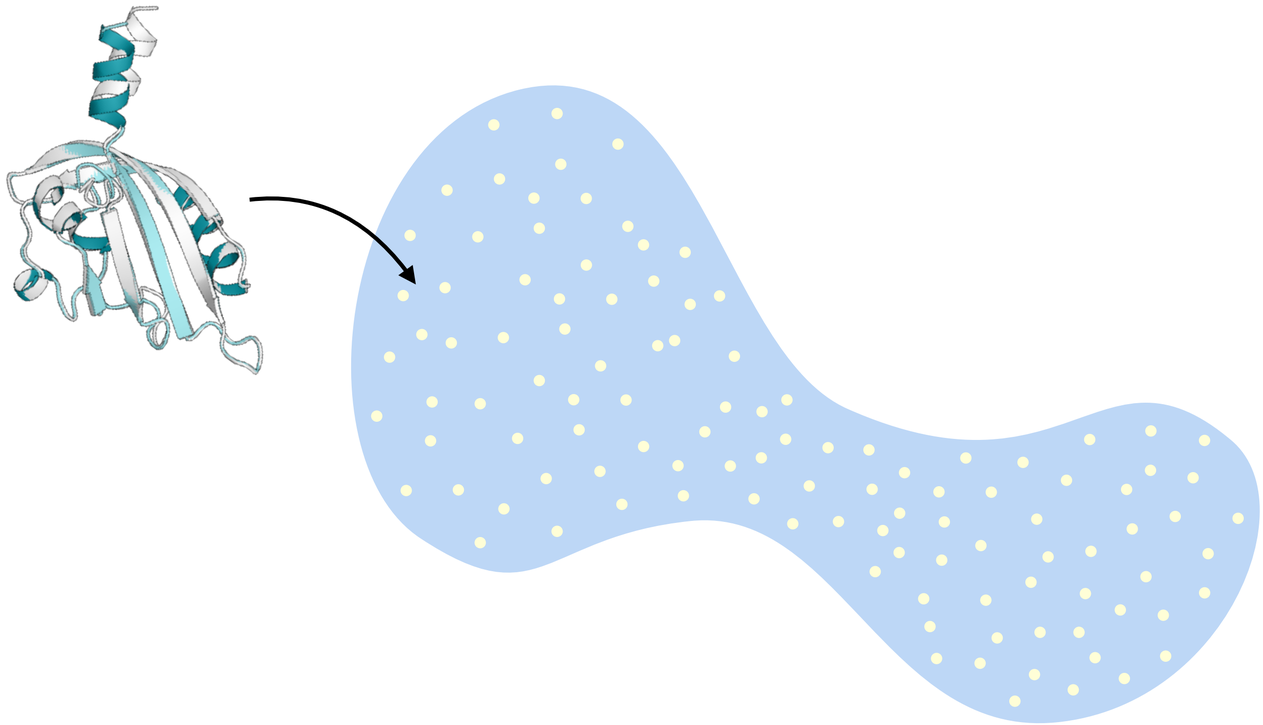

Data $D = \{x_i\}_{i=1}^{n}$

Generative Discovery Paradigm

valid design space Ω

generable set of pre-trained gen. model Ω_{pre}

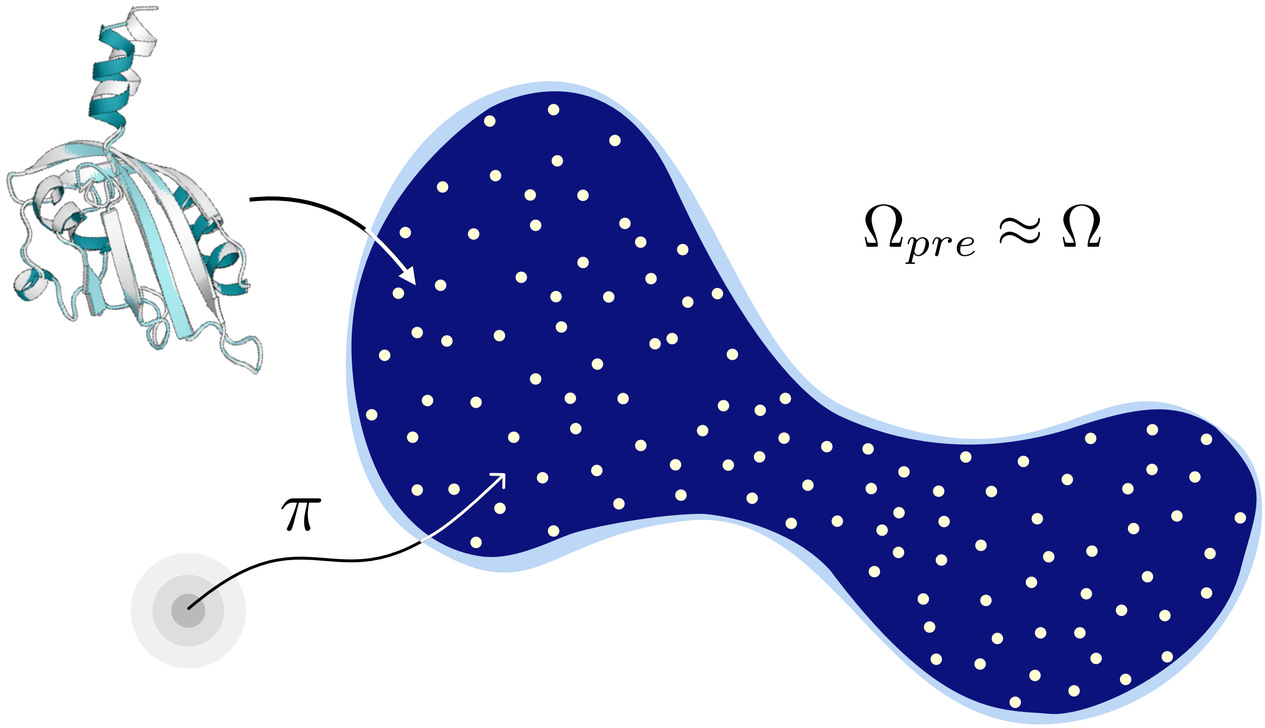

$\Omega_{pre} \approx \Omega$

Data $D = \{x_i\}_{i=1}^{n}$

↓

Pre-trained generative model π

Generative priors (e.g., diffusion) **model approximate, data-driven action sets**.

The Dream: Discovery via Generative Optimization

valid design space Ω

generable set of pre-trained gen. model Ω_{pre}

Discovery problem

$$\underset{x \in \Omega}{argmax}\ f(x)$$

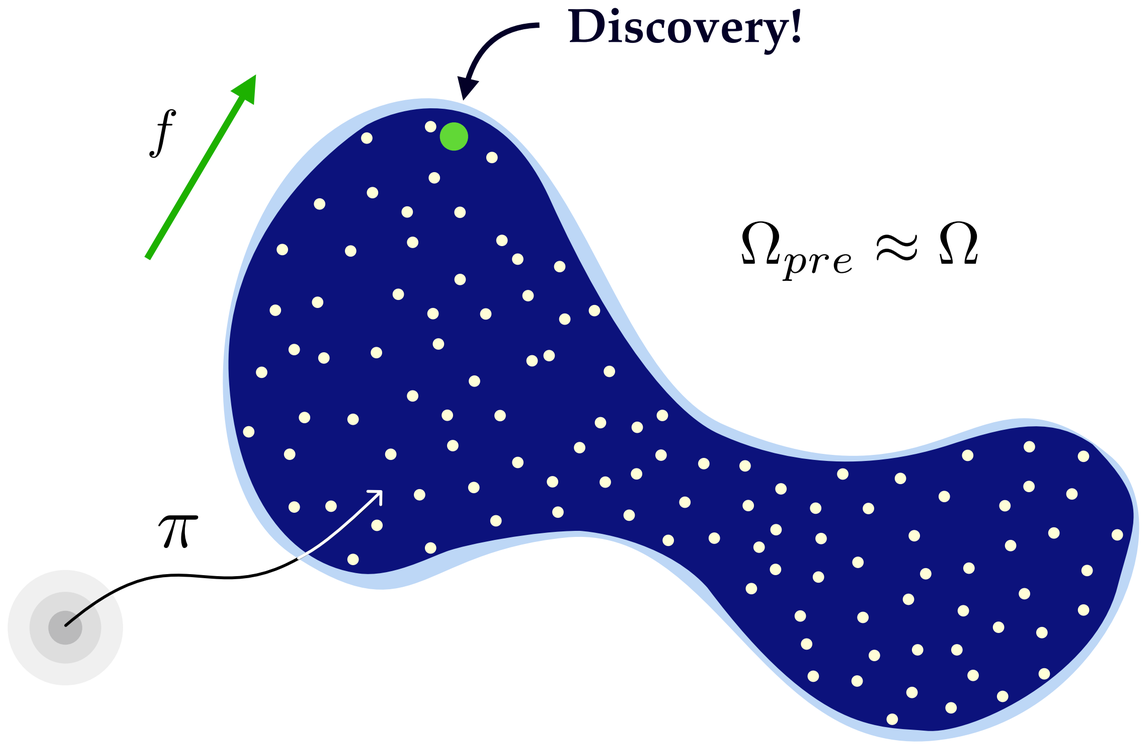

$\Omega_{pre} \approx \Omega$

Data $D = \{x_i\}_{i=1}^{n}$

↓

Pre-trained generative model π

↓

Adapt π to optimize f

(e.g., via Guidance, RL)

Generative priors (e.g., diffusion) **model approximate, data-driven action sets**.

Fundamental Challenge of Generative Discovery

valid design space Ω

generable set of pre-trained gen. model Ω_{pre}

Discovery problem

$$\underset{x \in \Omega}{argmax}\ f(x)$$

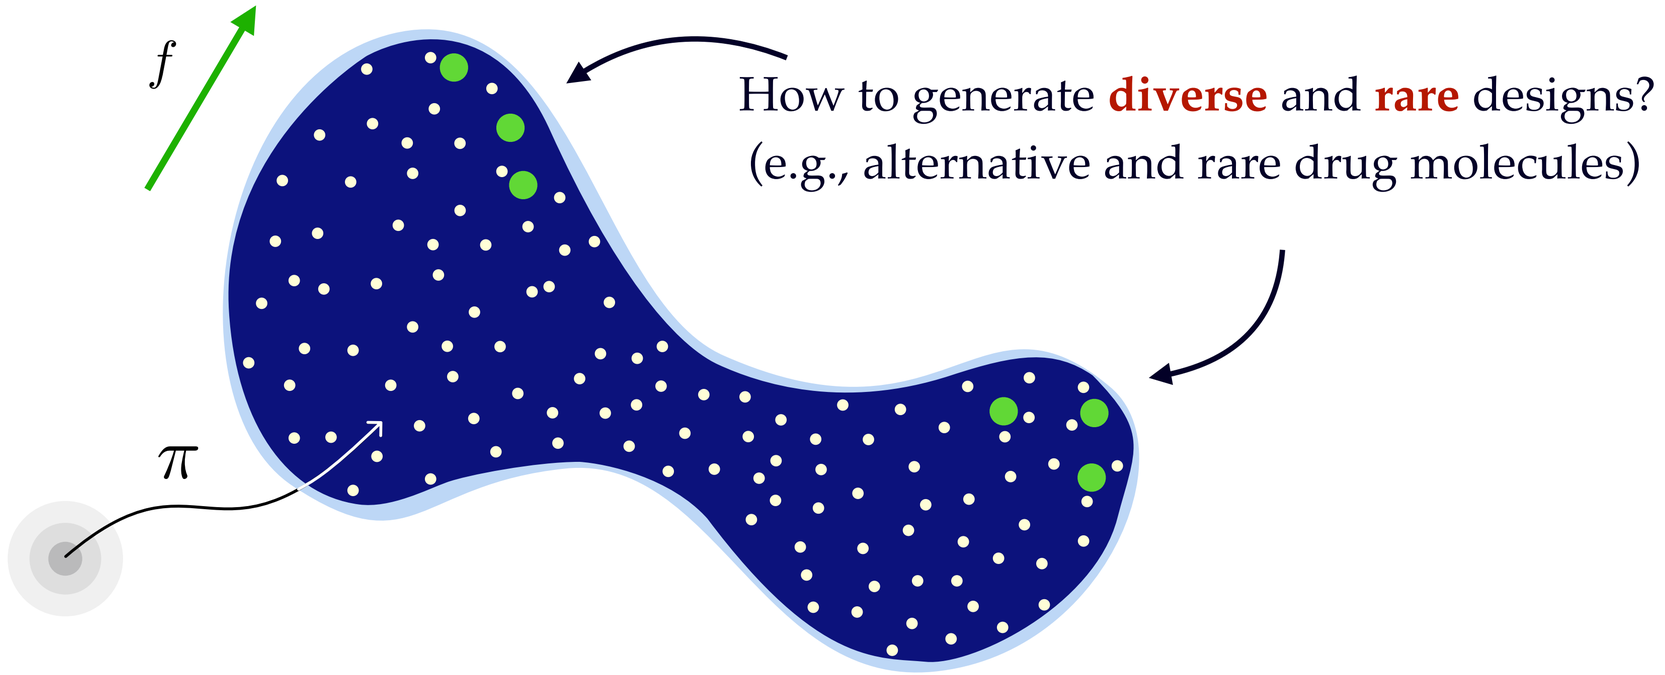

How to generate **diverse** and **rare** designs?
(e.g., alternative and rare drug molecules)

Fundamental Challenge of Generative Discovery

valid design space Ω

generable set of pre-trained gen. model Ω_{pre}

Discovery problem

$$\underset{x \in \Omega}{argmax}\ f(x)$$

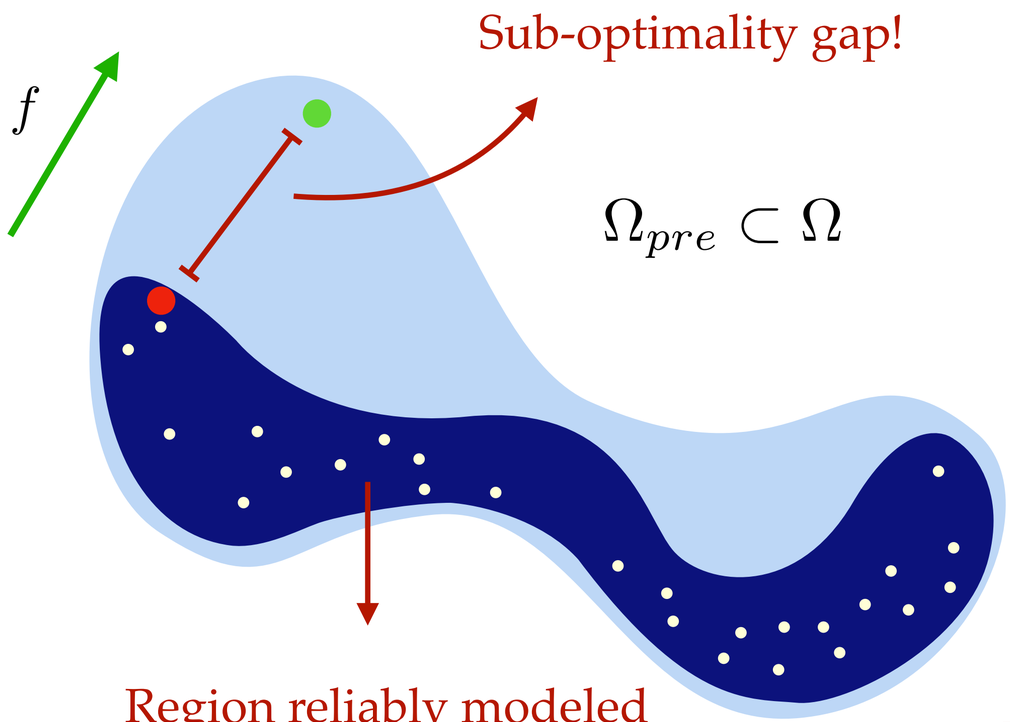

Region reliably modeled by generative priors is limited by training data.

Data is often scarce and with poor coverage in natural sciences. (e.g., enzymes, materials)

Fundamental Challenge of Generative Discovery

valid design space Ω

generable set of pre-trained gen. model Ω_{pre}

Discovery problem

$$\underset{x \in \Omega}{argmax}\ f(x)$$

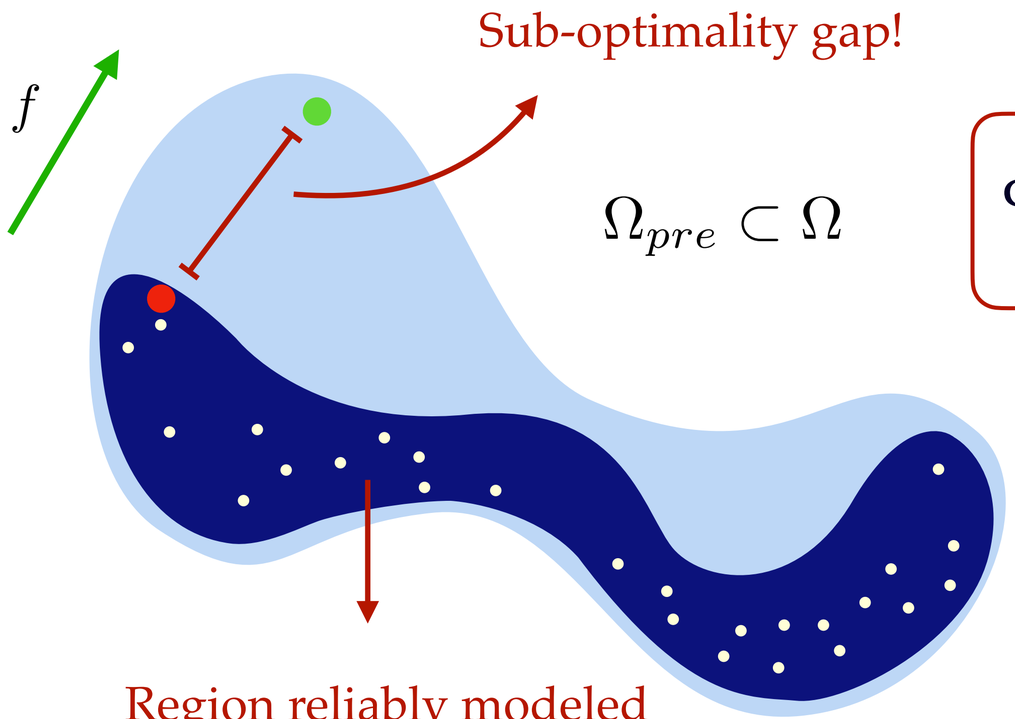

$\Omega_{pre} \subset \Omega$

Takeaway

Generative discovery requires the machine to **expand (learn) its action space.**

Region reliably modeled by generative priors is limited by training data.

Data is often scarce and with poor coverage in natural sciences. (e.g., enzymes, materials)

Fundamental Challenge of Generative Discovery

Discovery problem

$$\underset{x \in \Omega}{argmax}\ f(x)$$

High-level discovery goal

*"Discover novel **peptides** with promising **anticancer** activity"*

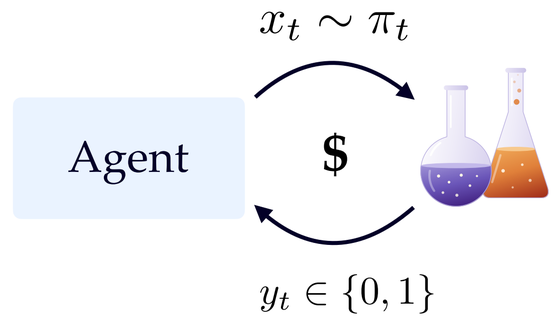

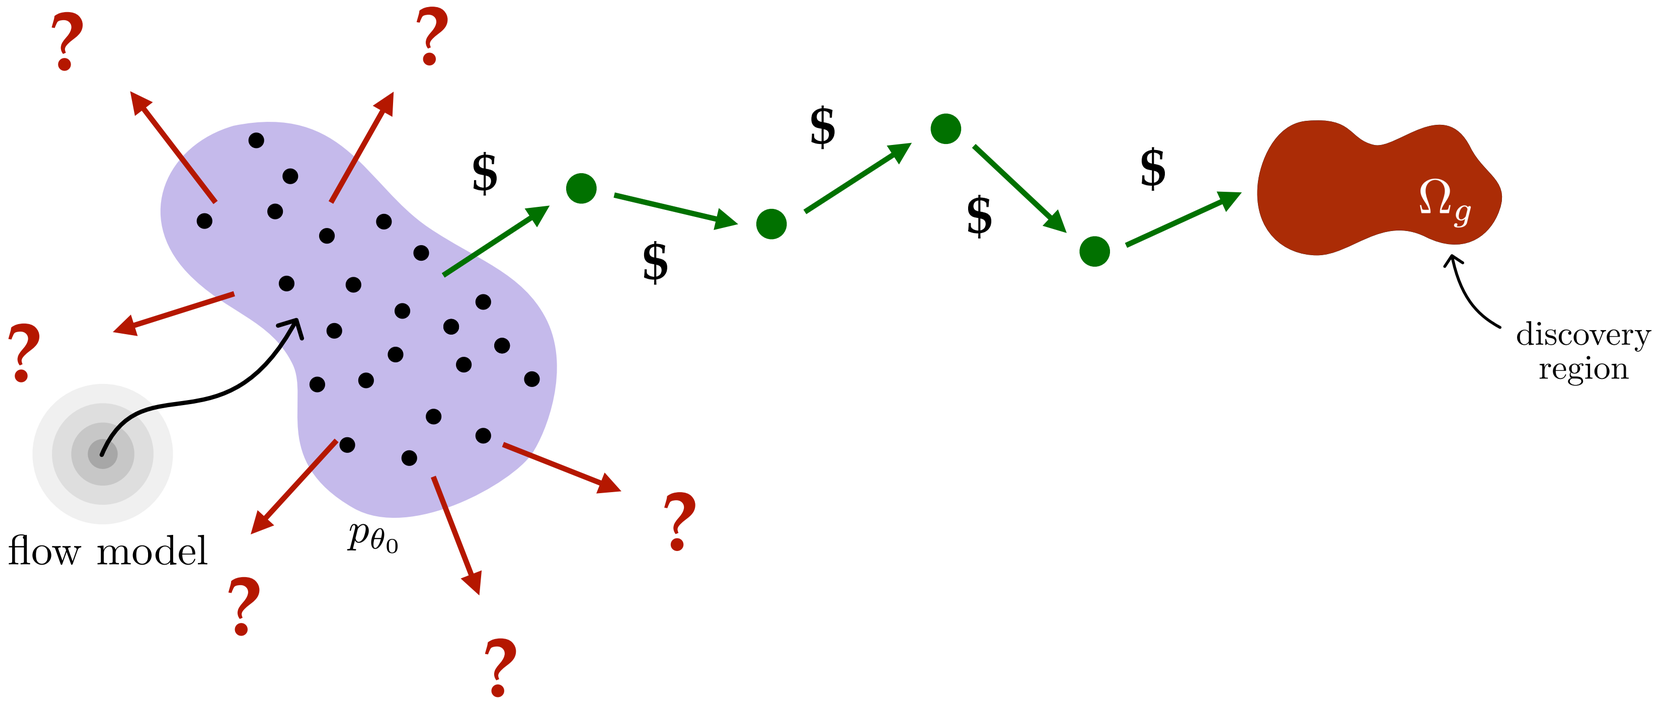

Generative Discovery

Learning to generate **diverse and rare** designs, **expanding beyond the prior**, toward **high-level discovery goals**.

Generative Discovery

Learning to generate **diverse and rare** designs, **expanding beyond the prior**, toward **high-level discovery goals**.

Distributional Flow Control for Discovery

Which candidates matter?

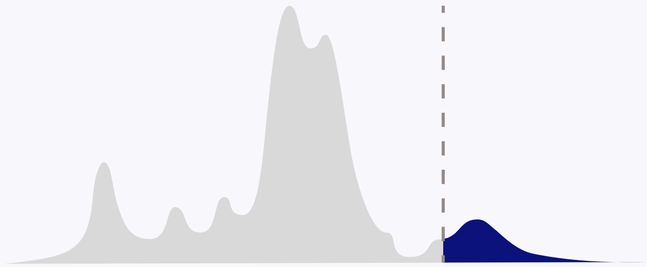

Out-of-Distribution Flow Modeling

Which regions can we reach?

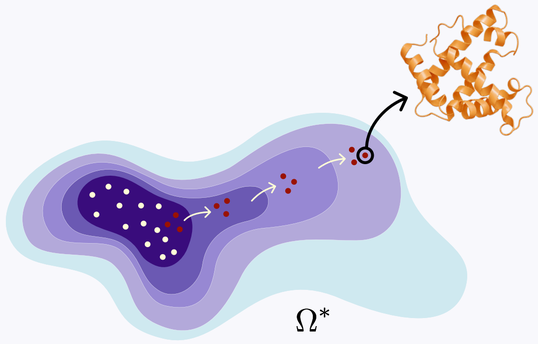

Sample-Efficient Discovery via LLM Distillation

What feedback guides us there?

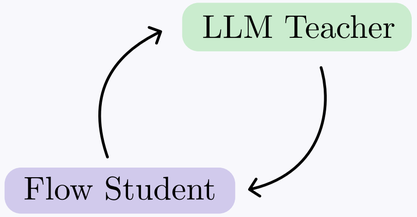

Distributional Flow Control for Discovery

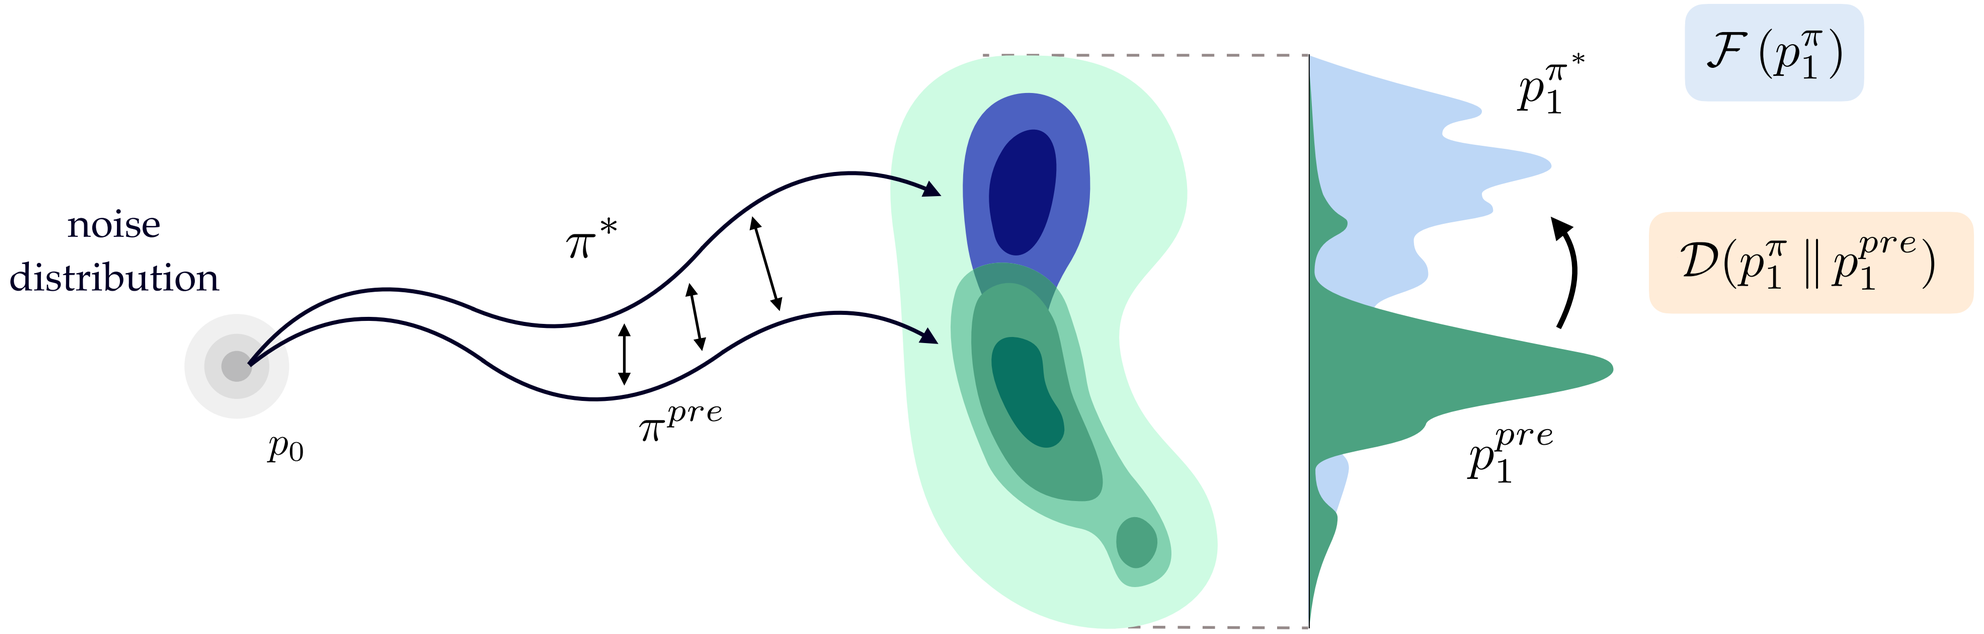

$\mathcal{F}$ **Discovery utility**

$\mathcal{D}$ **Divergence to preserve prior information**

Distributional Generative Control

$$\pi^* \in \operatorname*{argmax}_{\pi} \; \mathcal{F}\,(p_1^\pi) - \alpha \mathcal{D}(p_1^\pi \,\|\, p_1^{pre})$$

Distributional Generative Control

$$\pi^* \in \operatorname{argmax} \; \mathcal{F}(p_1^\pi) - \alpha \mathcal{D}(p_1^\pi \,\|\, p_1^{pre})$$

Extreme event generation

$$\mathrm{CVaR}^f_{1-\beta}(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \geq \mathrm{VaR}^f_{1-\beta}(p_1^\pi)]$$

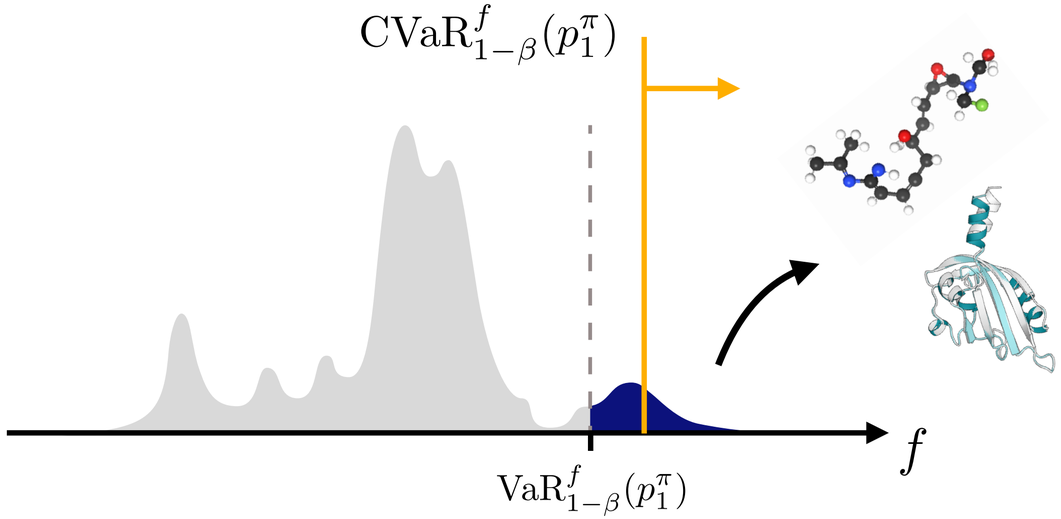

Distributional control is essential for discovery (e.g., diversity, extreme, rare events etc.)

Distributional Generative Control

$$\pi^* \in \operatorname{argmax} \; \mathcal{F}(p_1^\pi) - \alpha \mathcal{D}(p_1^\pi \,\|\, p_1^{pre})$$

Extreme event generation

$$\mathrm{CVaR}^f_{1-\beta}(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \geq \mathrm{VaR}^f_{1-\beta}(p_1^\pi)]$$

Diverse generation

$$\mathcal{H}(p_1^\pi) := -\mathbb{E}_{x\sim p^\pi}[\log p^\pi(x)]$$

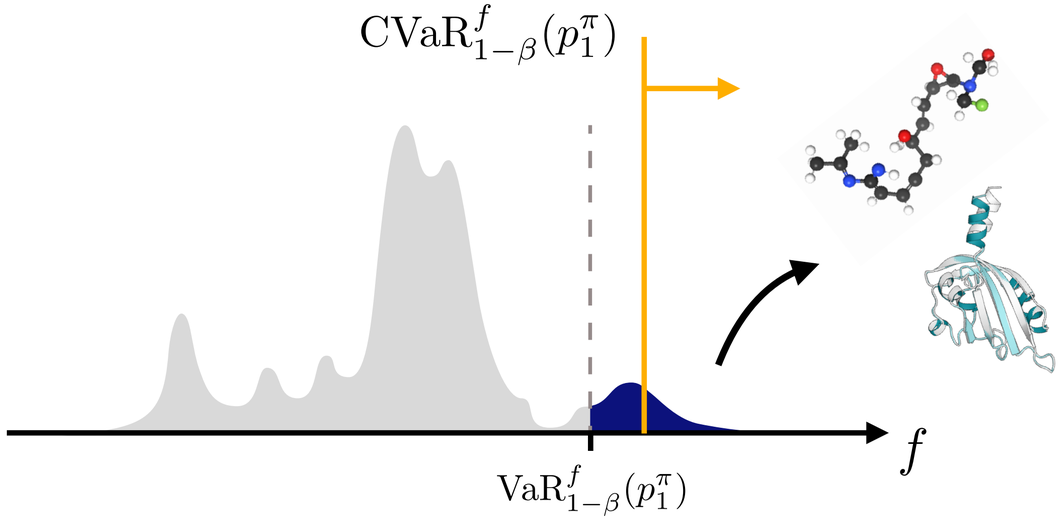

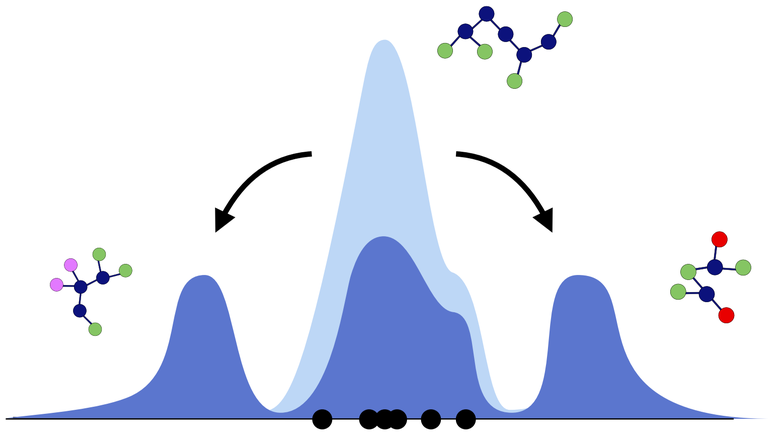

Distributional control is essential for discovery (e.g., diversity, extreme, rare events etc.)

Extreme Event Generation via CVaR

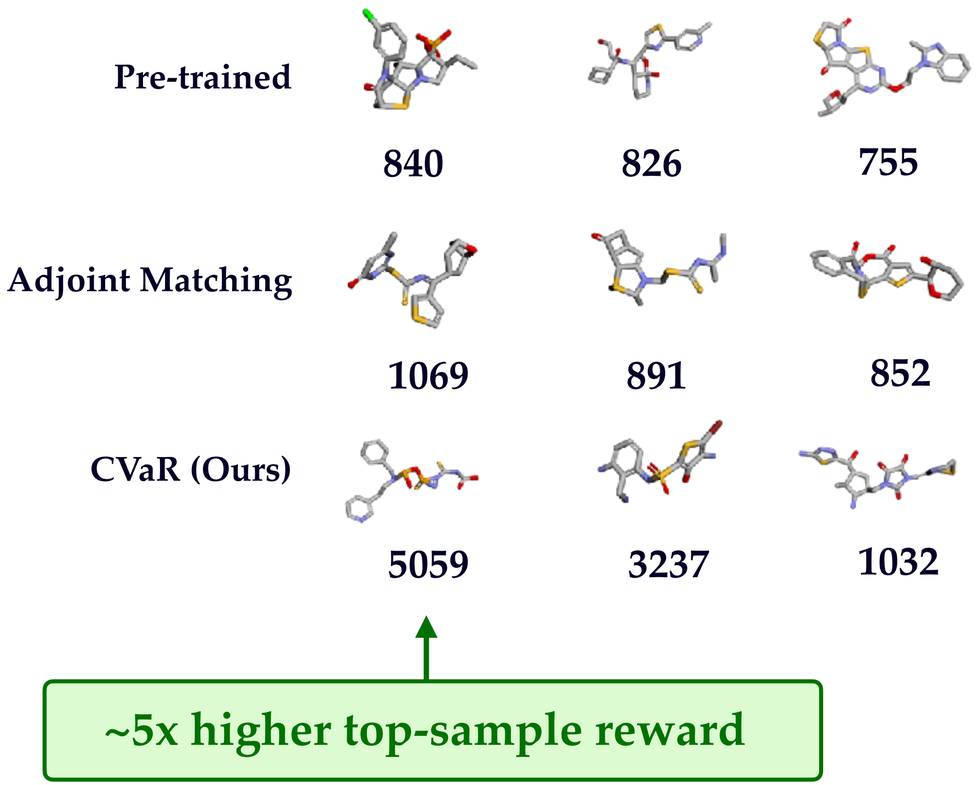

FDC significantly **improves top-samples**

[Dunn and Koes, 2024] [Axelrod and Bombarelli, 2022]

Reward = Negative energy via GFN1-xTB, Validity = RDKit Sanitizer

Extreme Event Generation via CVaR

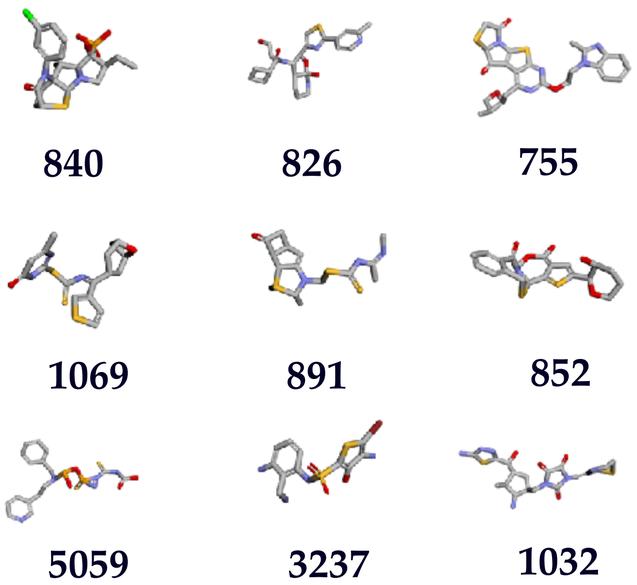

~5x higher top-sample reward

FDC significantly **improves top-samples**

Diverse Generation via Entropy

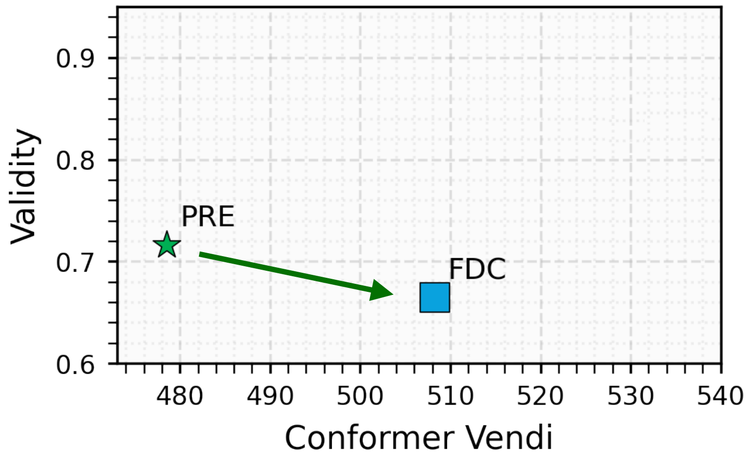

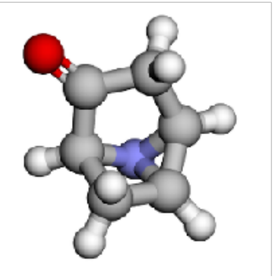

FDC increases structural diversity
while maintaining **high validity**

[Dunn and Koes, 2024] [Axelrod and Bombarelli, 2022]

Reward = Negative energy via GFN1-xTB, Validity = RDKit Sanitizer

Out-of-Distribution Generative Modeling

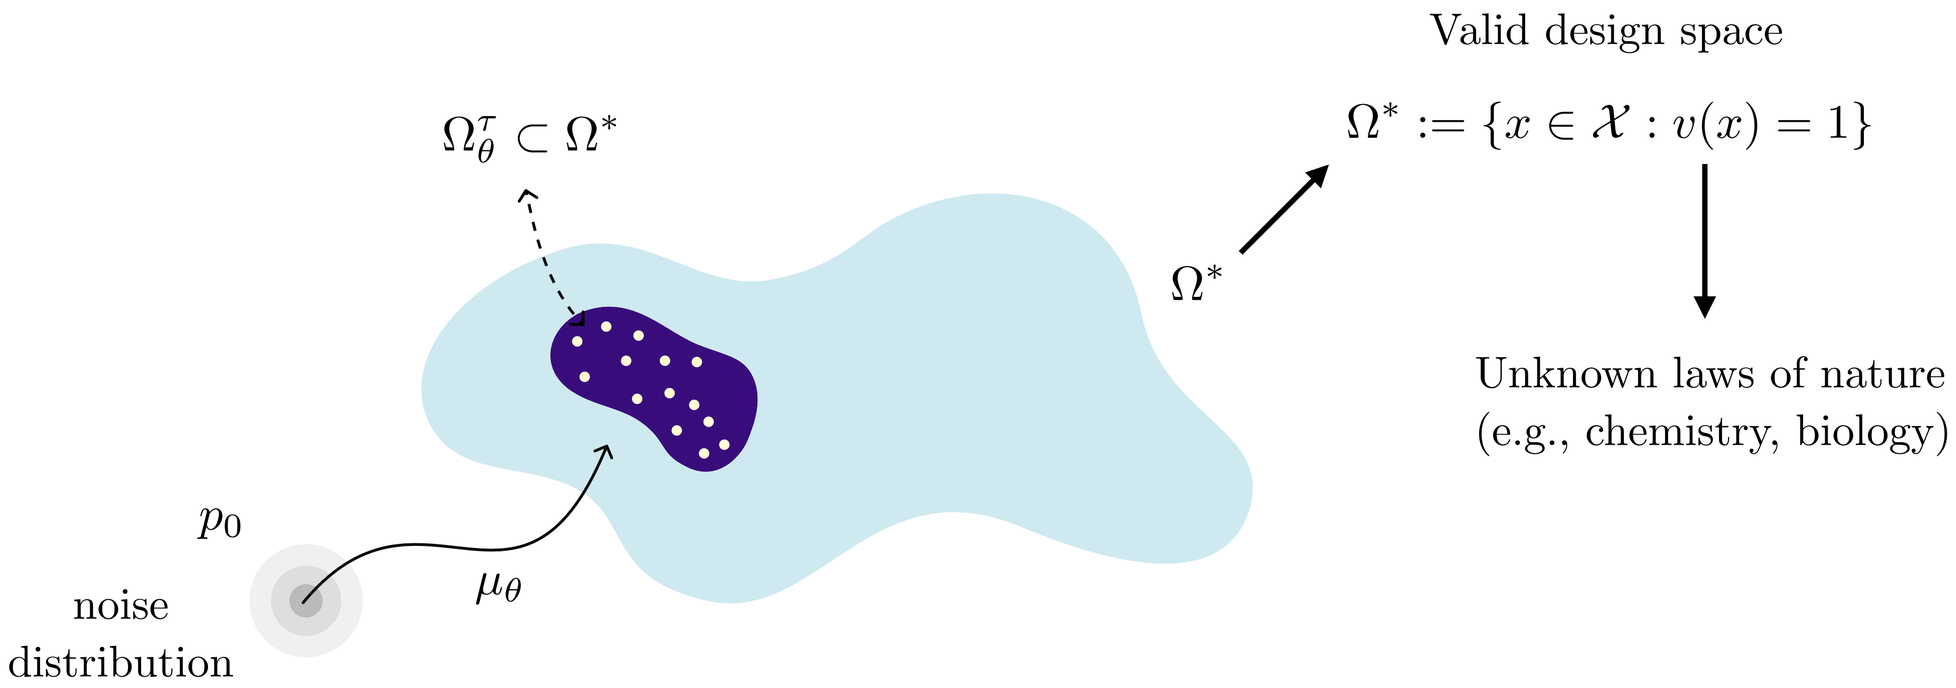

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

$$\Omega_\theta^\tau := \{x \in \mathcal{X} : p_1^\theta(x) \geq \tau\}$$

→ **generable set $\approx$ finite-sample coverage**

Generable Set Expansion: Beyond Distribution Matching

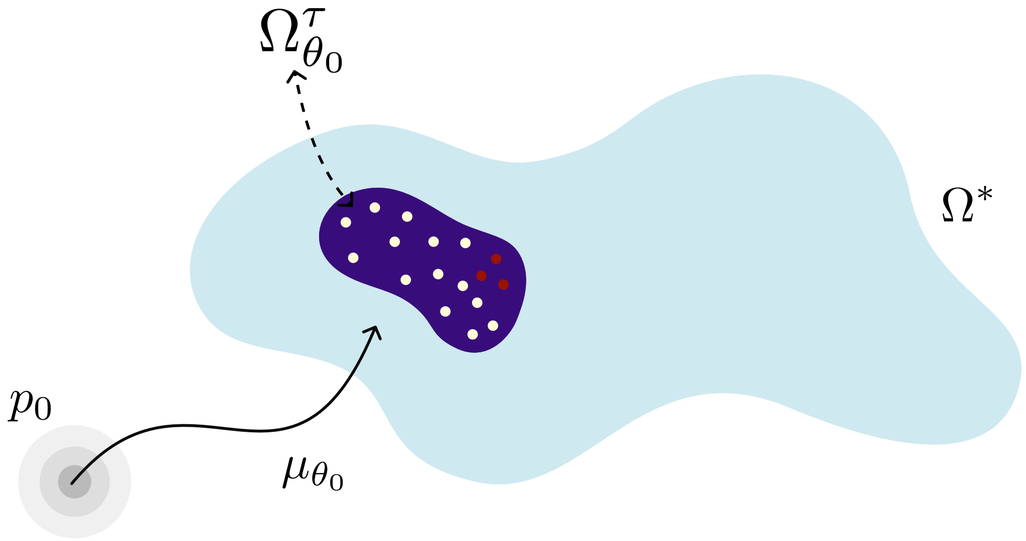

Black-box noisy binary feedback

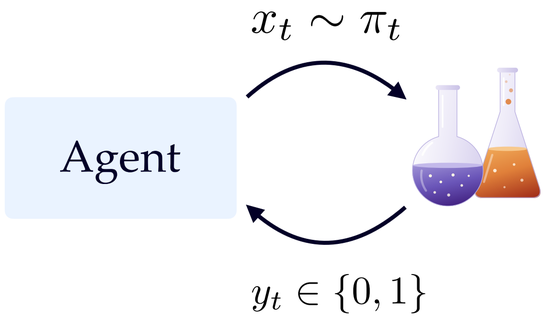

Generable Set Expansion: Beyond Distribution Matching

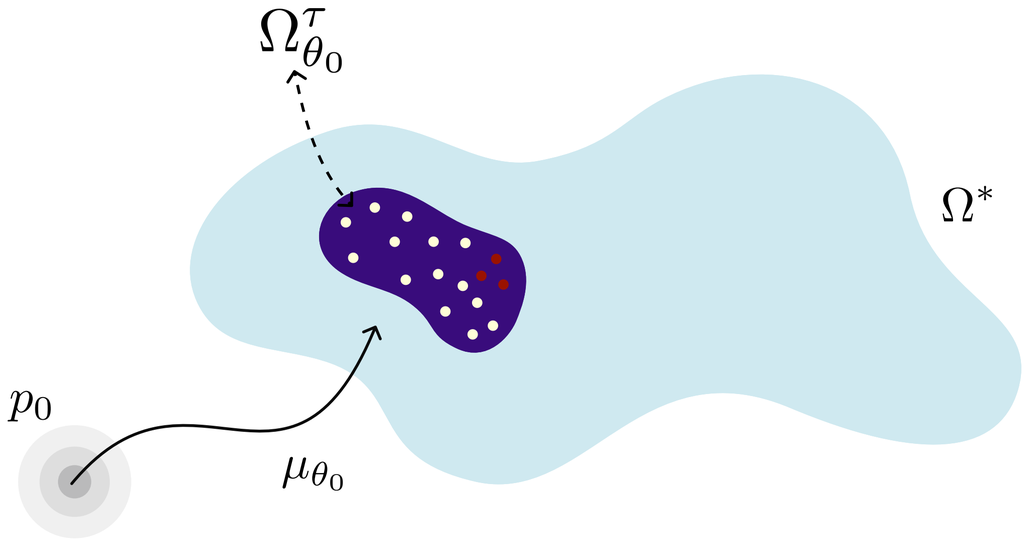

Black-box noisy binary feedback

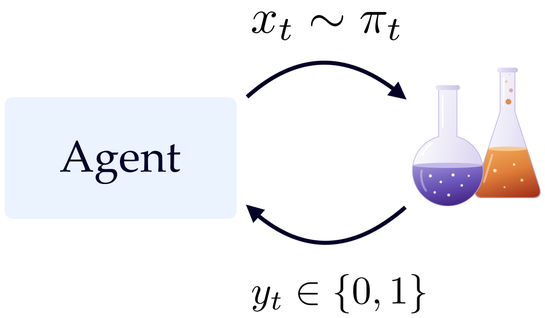

$\theta^{\text{pre}} = \theta_0$ $\Omega^{\tau}_{\theta_0}$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

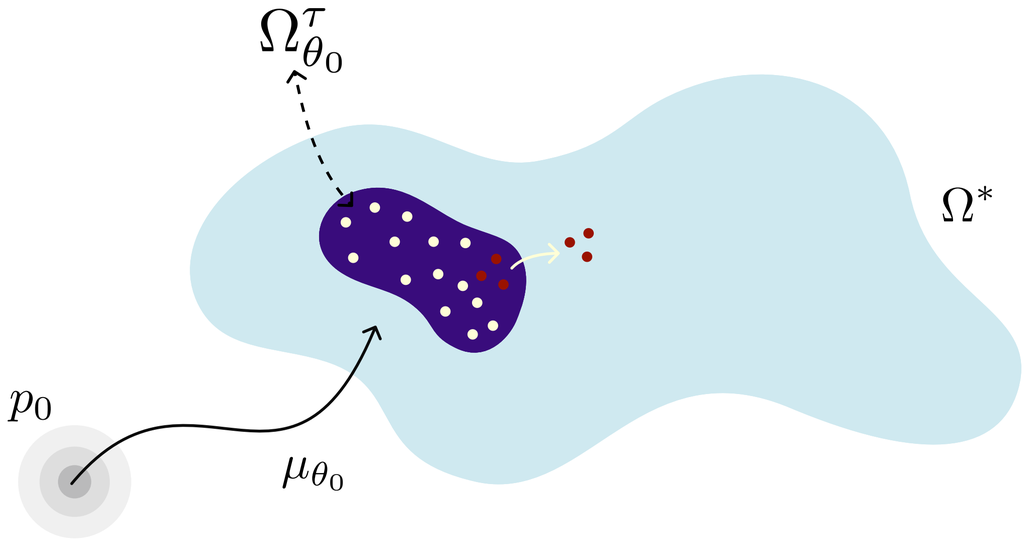

Black-box noisy binary feedback

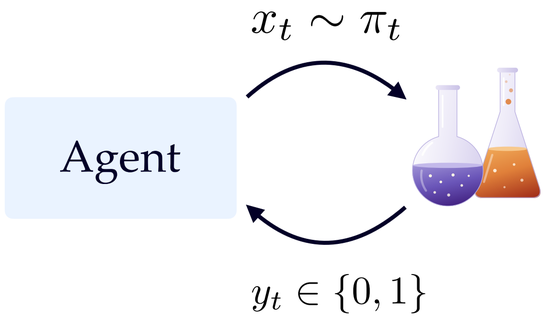

$\theta^{\text{pre}} = \theta_0$ $\Omega^{\tau}_{\theta_0}$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

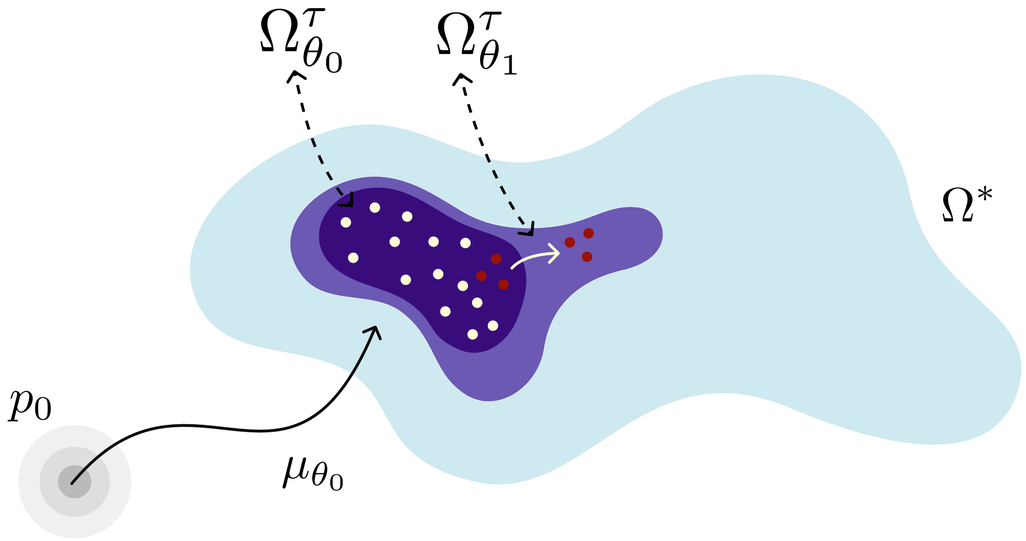

Black-box noisy binary feedback

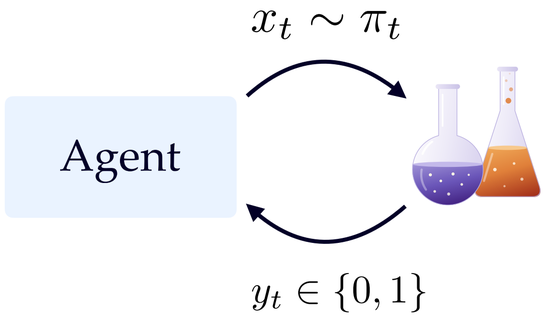

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1}$$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

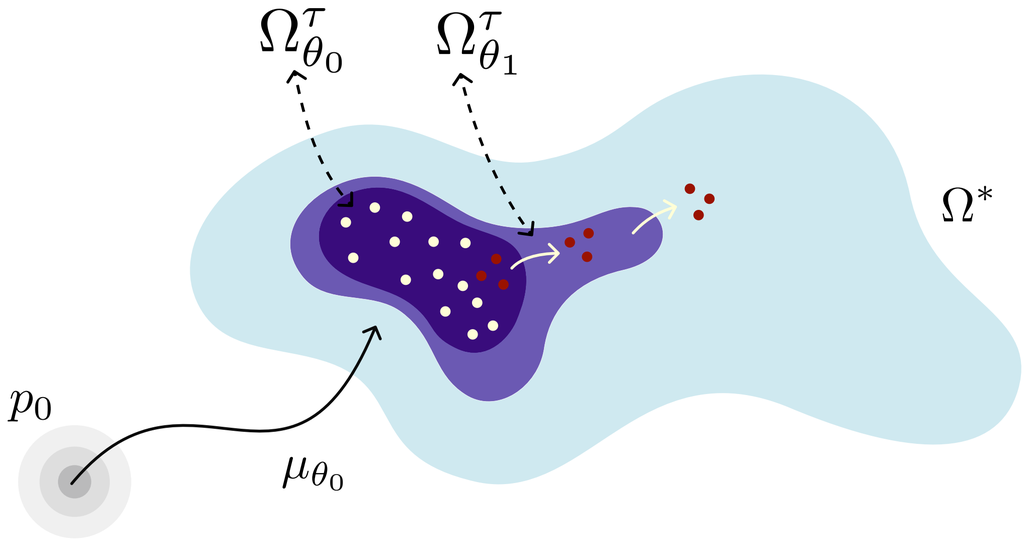

Black-box noisy binary feedback

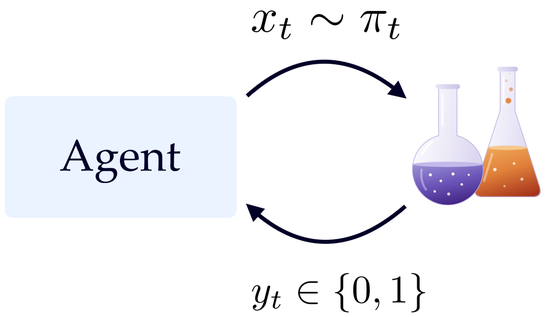

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq$$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

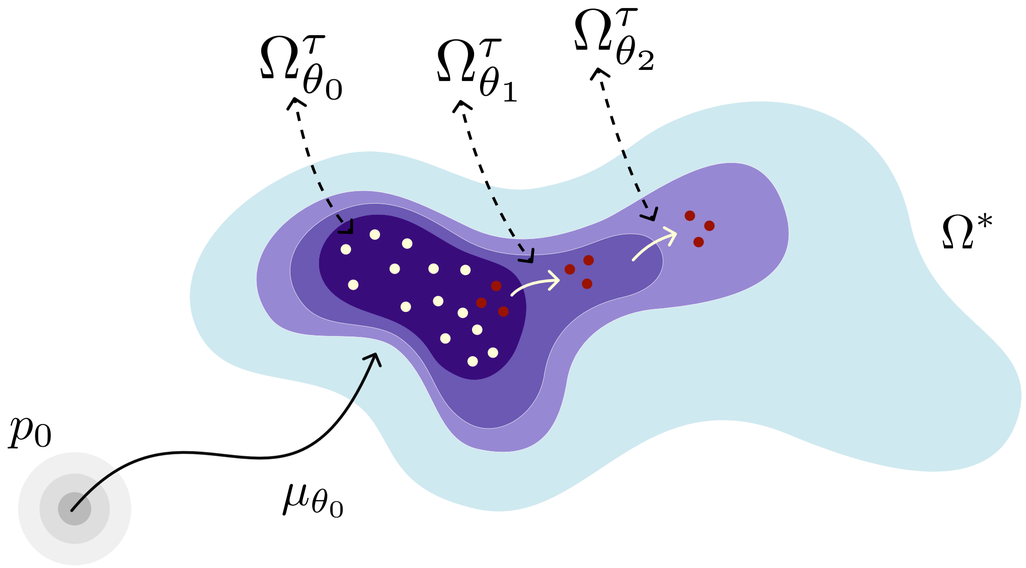

Black-box noisy binary feedback

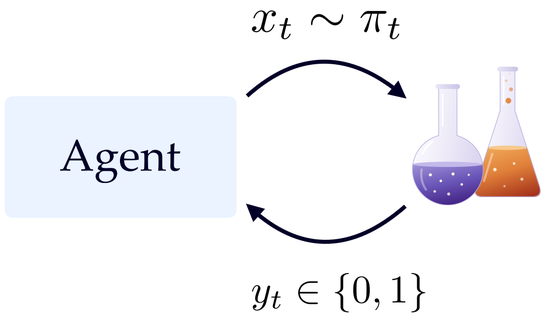

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots$$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

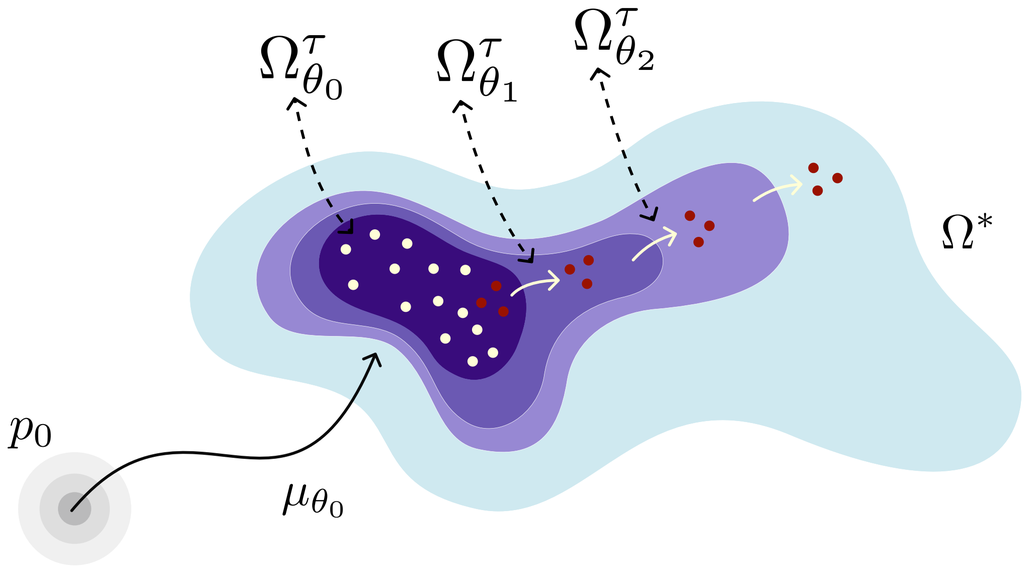

Black-box noisy binary feedback

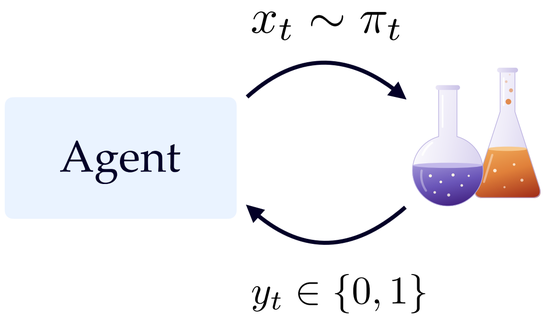

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq$$

Continual expansion via adaptation on valid **frontier samples**

Generable Set Expansion: Beyond Distribution Matching

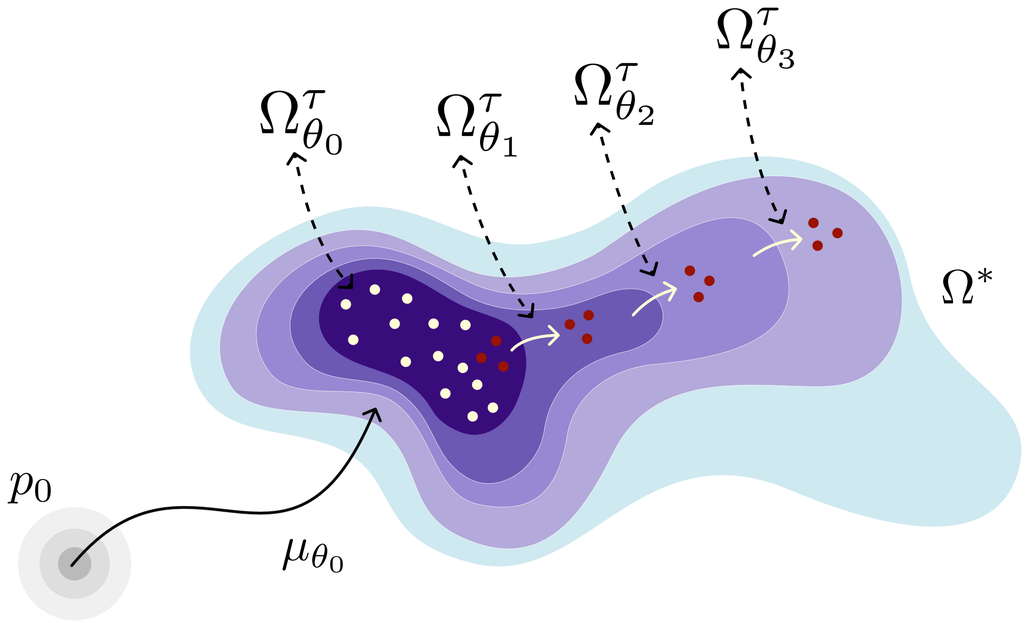

Black-box noisy binary feedback

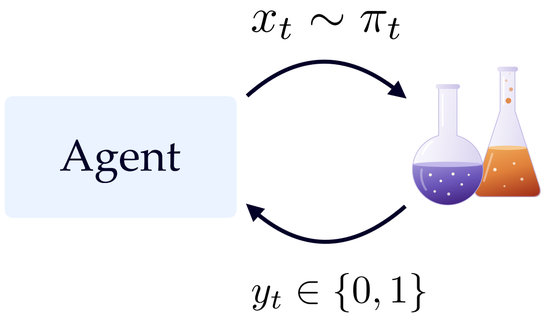

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

Continual expansion via adaptation on valid **frontier samples**

How to self-generate frontier samples and avoid model collapse?

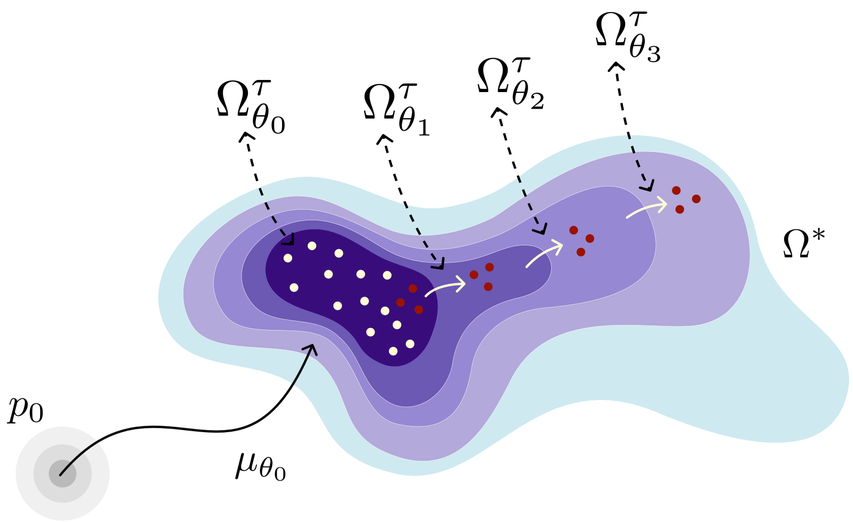

verifier uncertainty estimate

prior regularization

next data point

$$x_{t+1} \sim \tilde{p}_t \in \arg\max_q \mathbb{E}_{x\sim q}\left[\sigma_t\left(\phi_s^t(x)\right)\right] - \beta\,\mathrm{KL}\left(q \,\|\, p_1^{\theta_t}\right)$$

Inference-time flow steering for **active data collection.**

ActFlow with Chemical and Biological Data

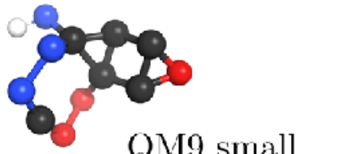

QM9 small
organic molecule

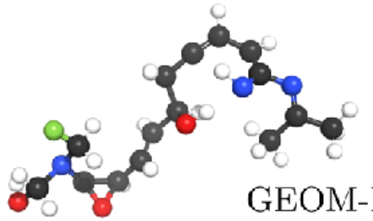

GEOM-Drugs
drug-like molecule

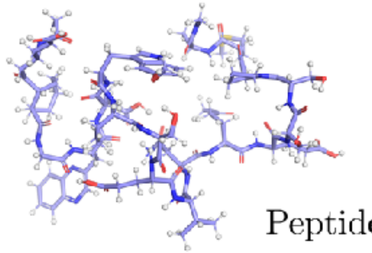

Peptide

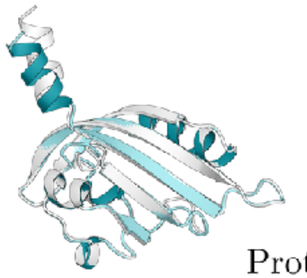

Protein

ActFlow for Small Organic and Drug-like Molecules

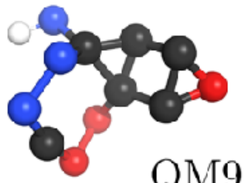

QM9 small
organic molecule

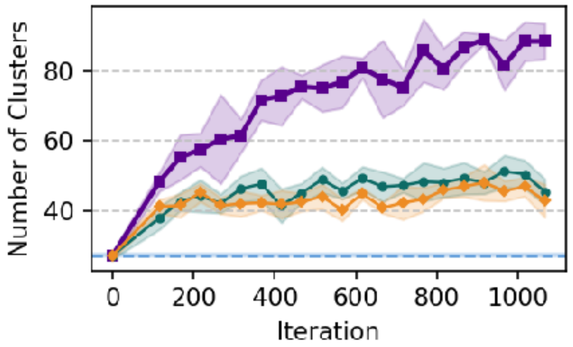

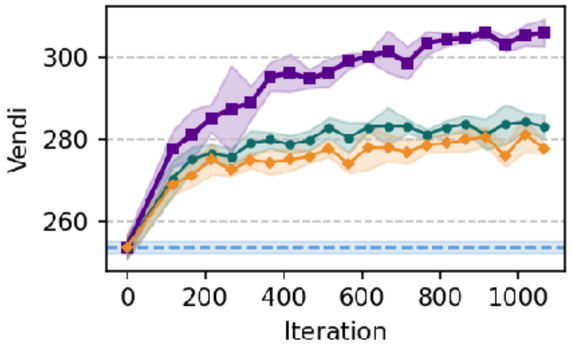

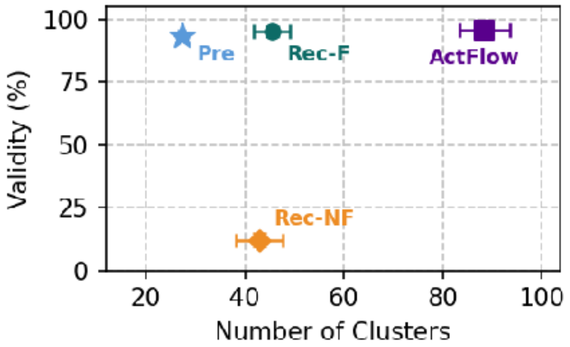

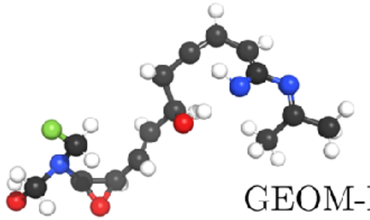

GEOM-Drugs
drug-like molecule

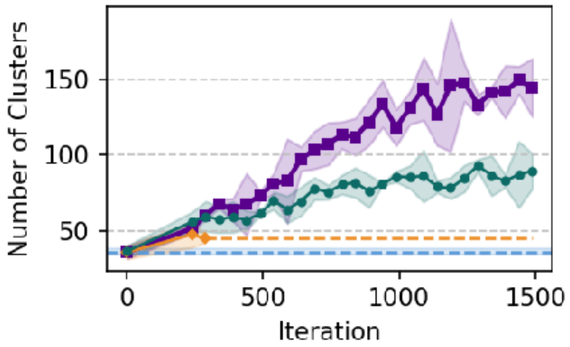

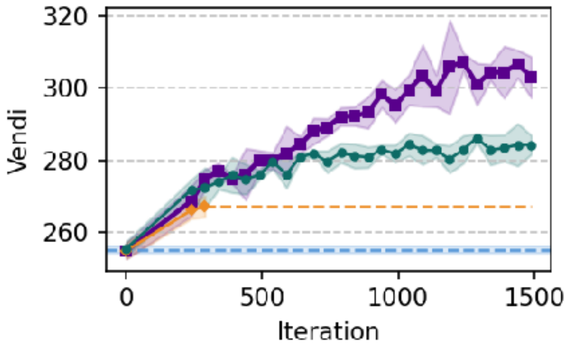

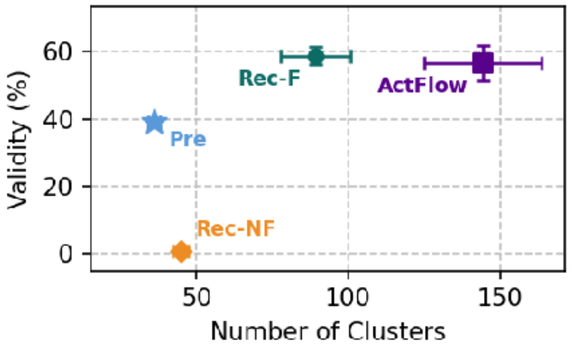

ActFlow on Proteins

Model collapse!

ActFlow self-expands on valid regions with **no model collapse**

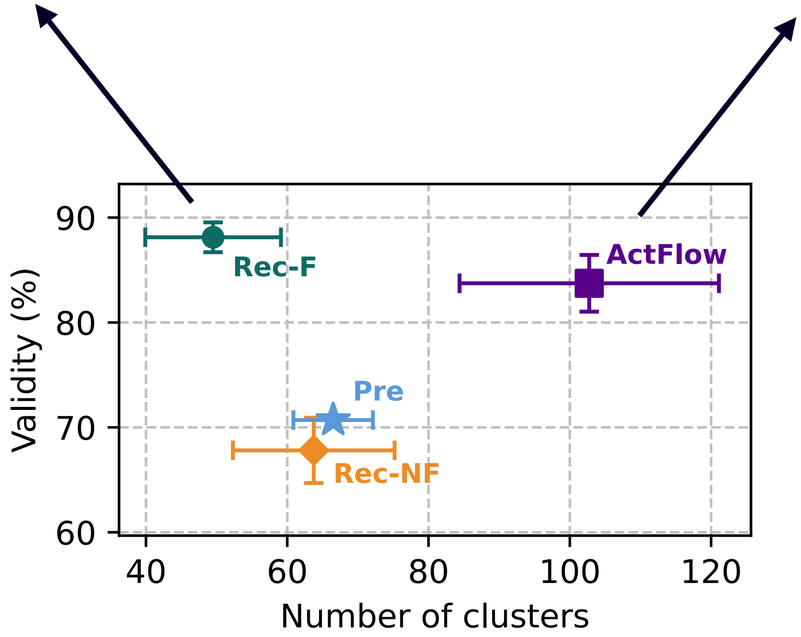

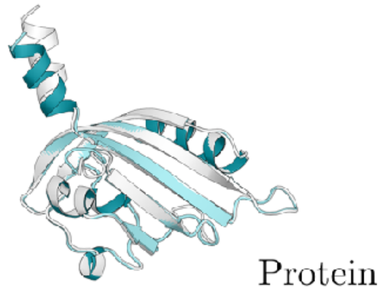

Sample-Efficient Goal-Directed Expansion for Discovery?

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'* $\longrightarrow$

Discovery region

$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$

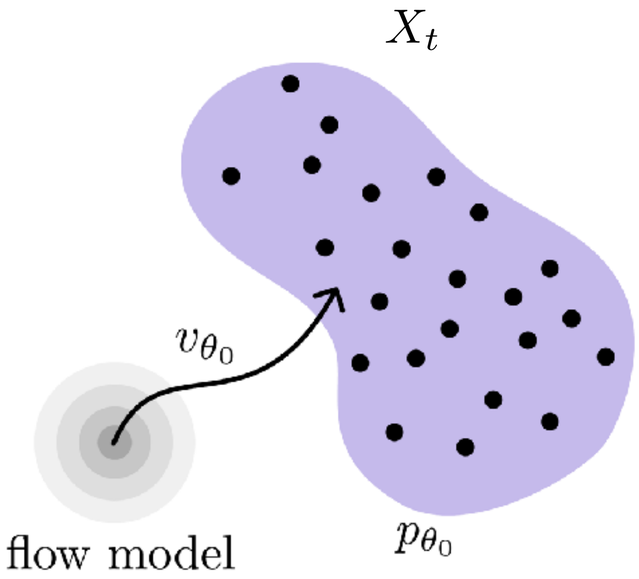

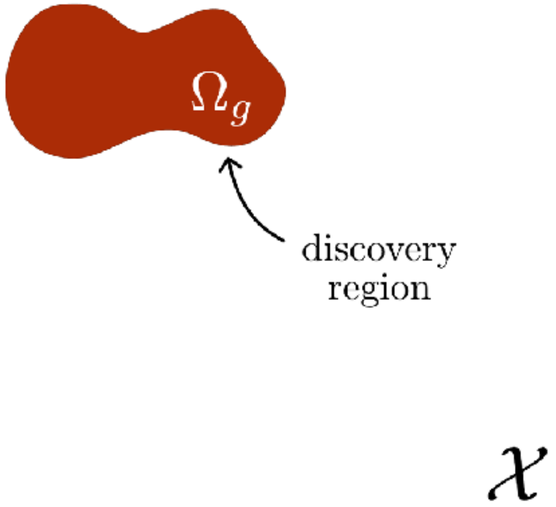

LLM-to-Flow Distillation (LFD)

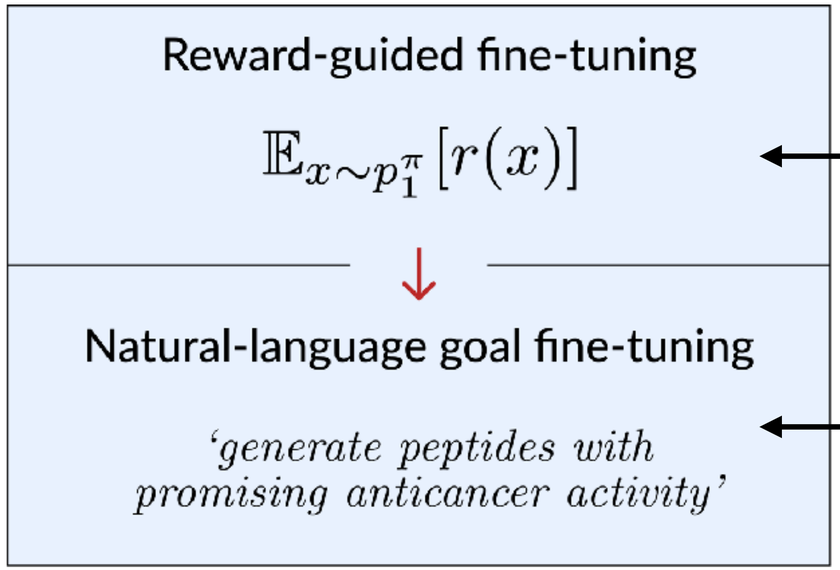

DPO, Flow-GRPO, Adjoint Matching, DiffusionNFT etc.

LFD (ours) (first method)

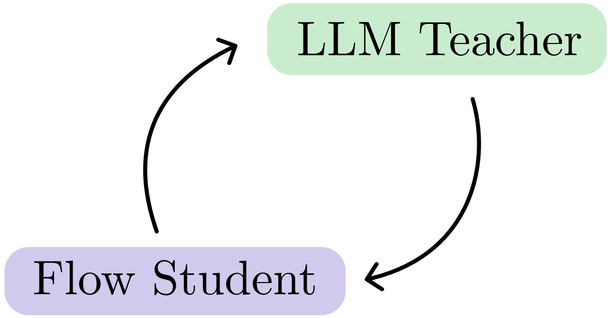

LLM-to-Flow Distillation (LFD) is the **first method** to **adapt any scientific generator to natural-language discovery goals**.

LFD requires **only LLM API-access** → Can use any **closed frontier model**

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

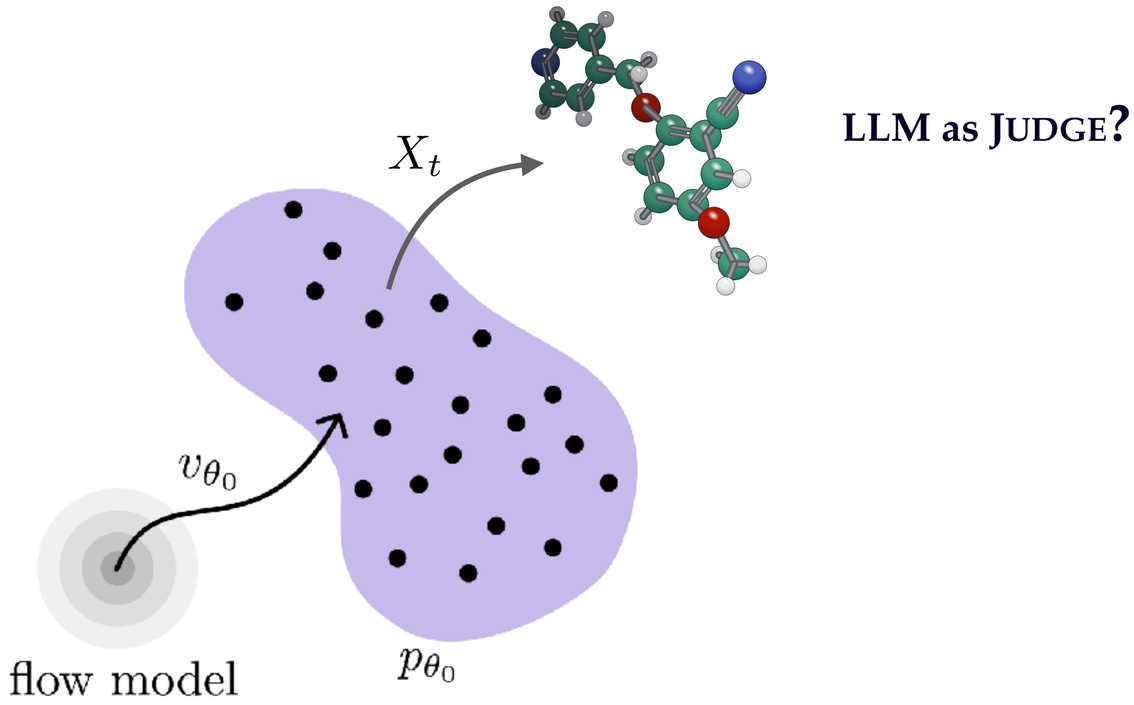

LLM as JUDGE?

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

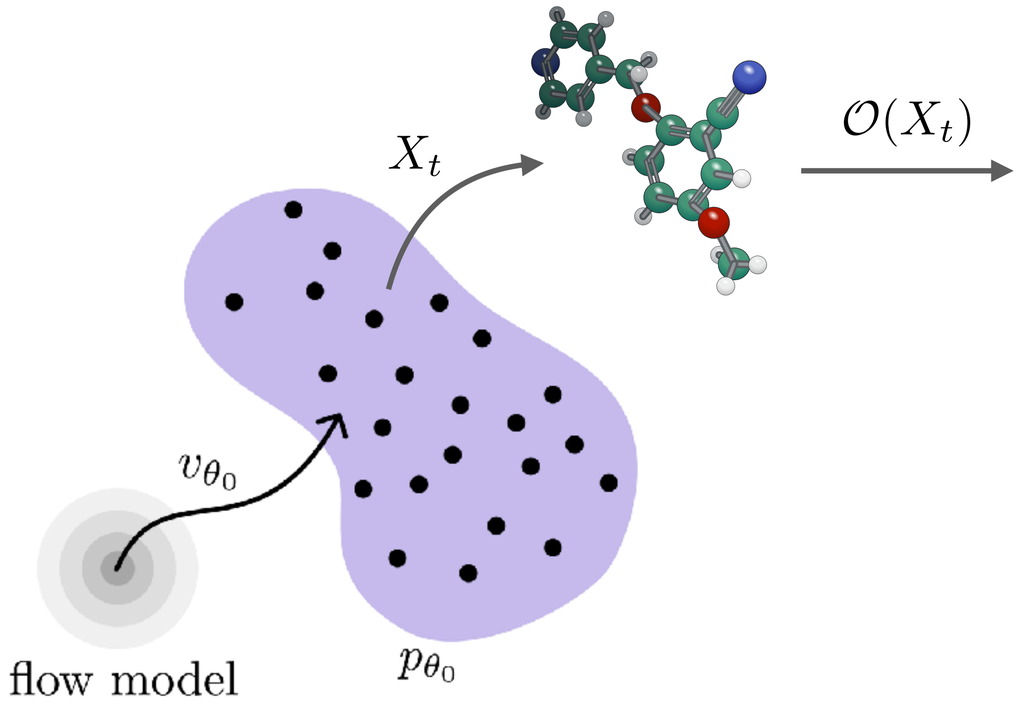

Observables

- SMILES (string)
- Ring counts
- Bond angles
- Atom types

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

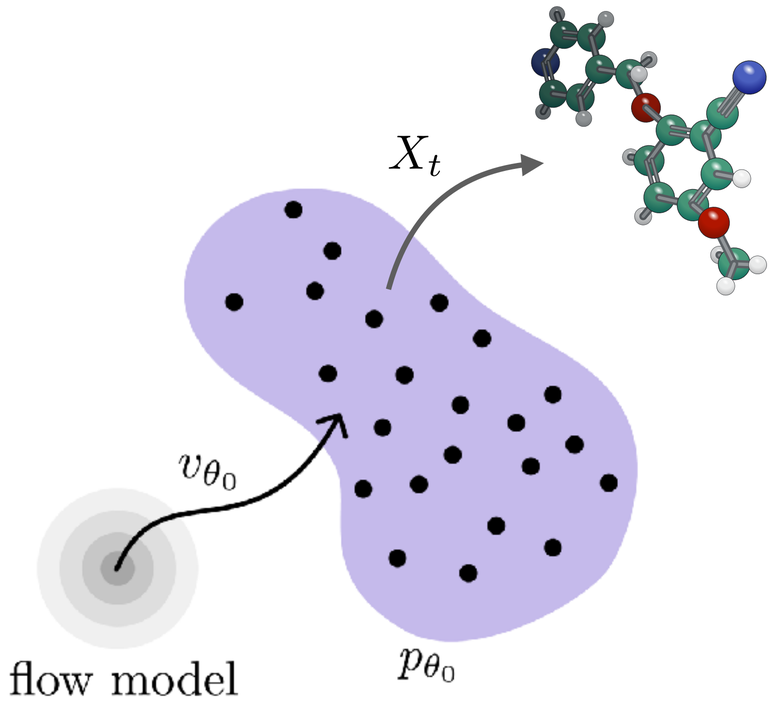

LLM as JUDGE?

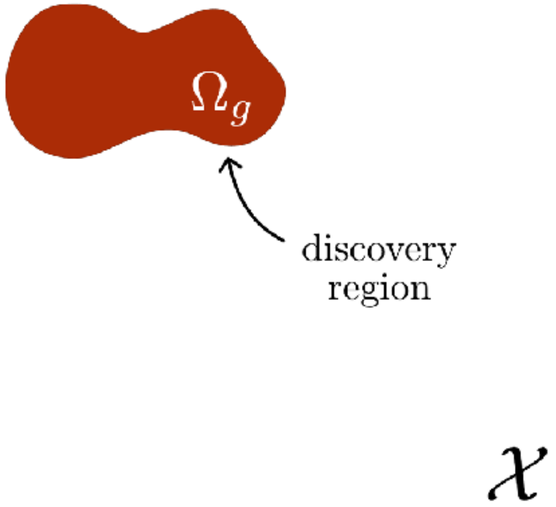

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$

$g_0 =$ *'prefer peptides that possess a net positive charge'*

$g_t \leftarrow \textsc{SetSubgoal}(g, g_{t-1}, \mathcal{O}(X_t))$

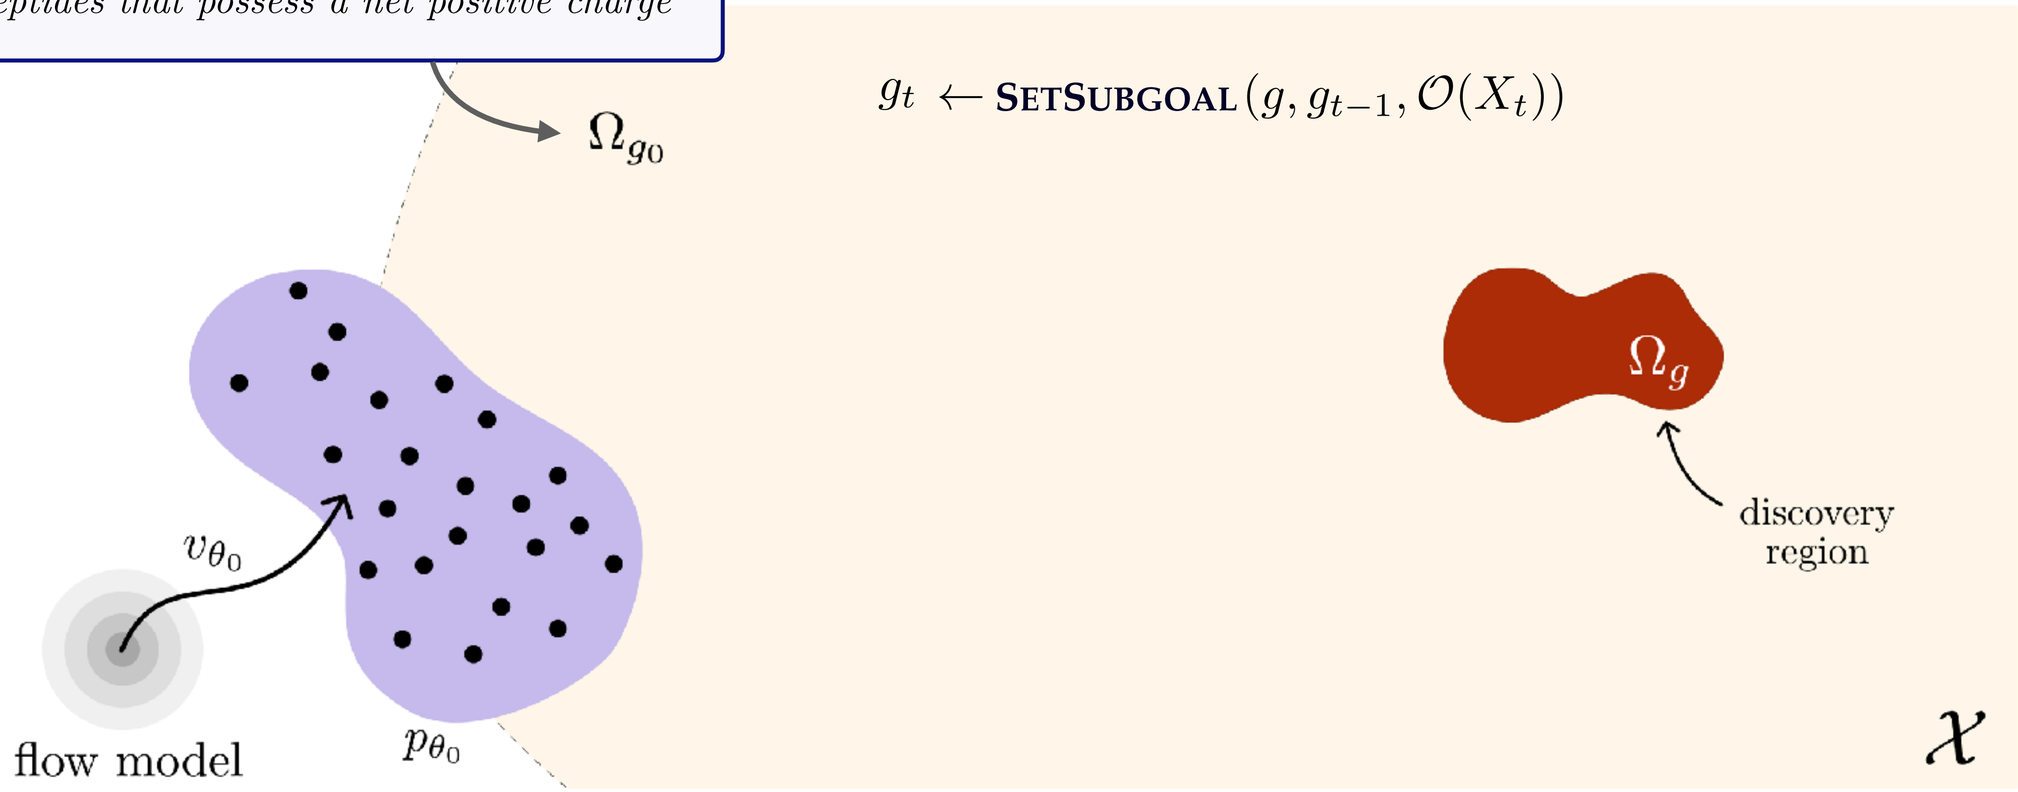

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$

$g_0 =$ *'prefer peptides that possess a net positive charge'*

$g_t \leftarrow \textbf{SETSUBGOAL}(g, g_{t-1}, \mathcal{O}(X_t))$

$D_t \leftarrow \textbf{LLM-JUDGE}\ (X_t, g_t)$

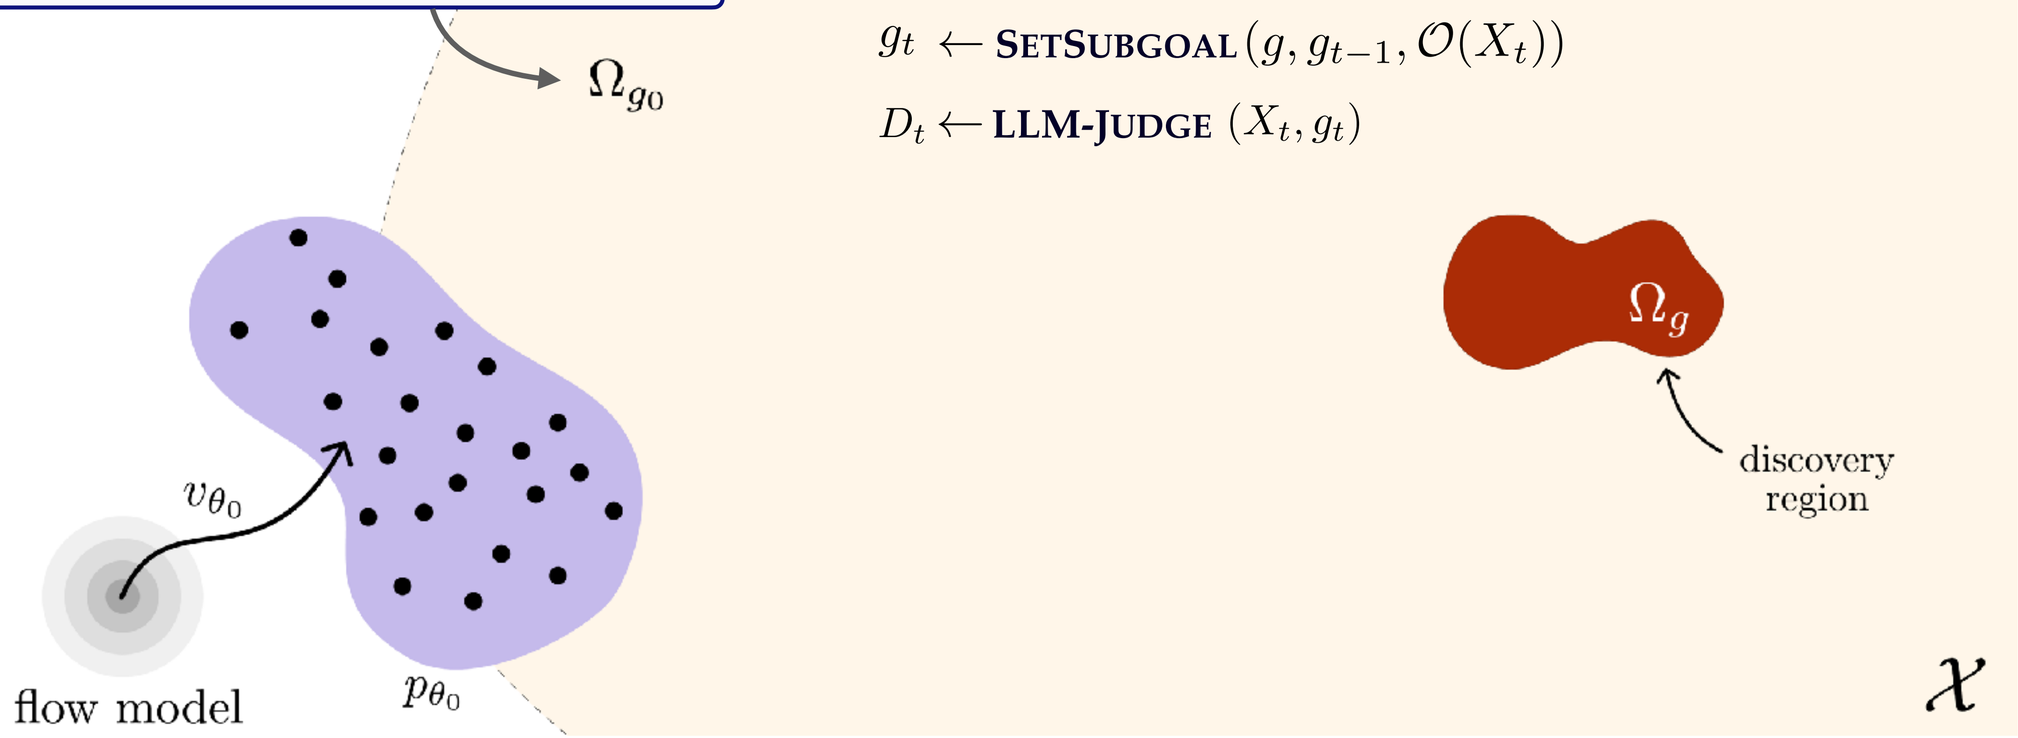

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$

$g_0 =$ *'prefer peptides that possess a net positive charge'*

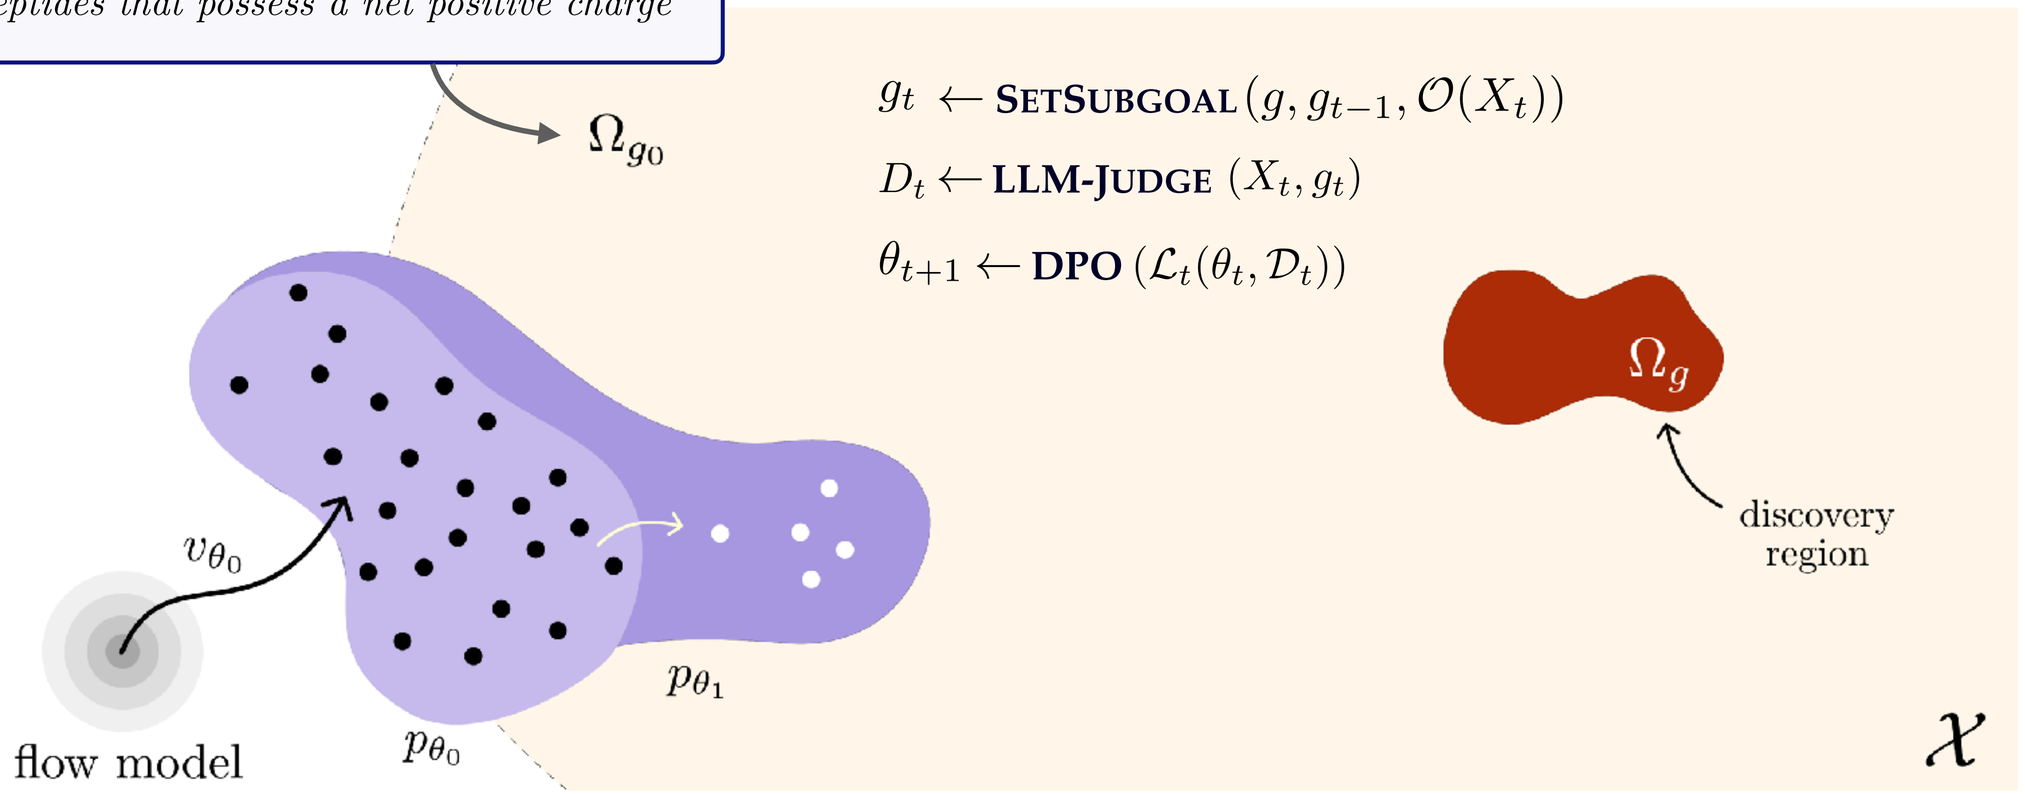

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

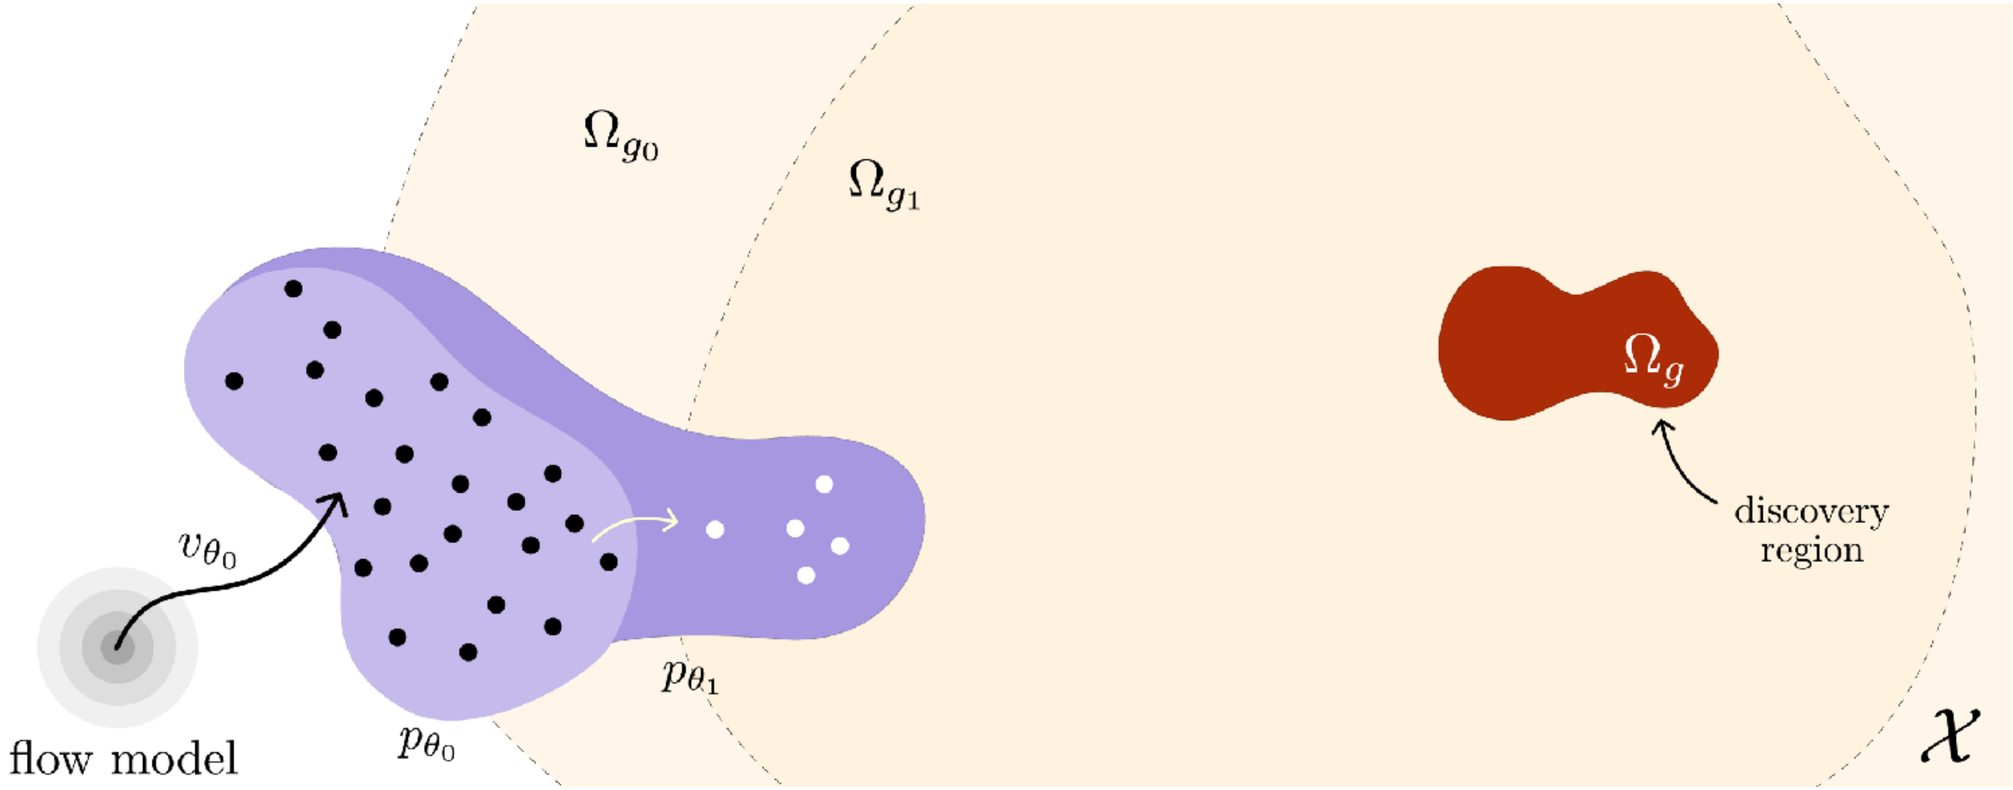

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

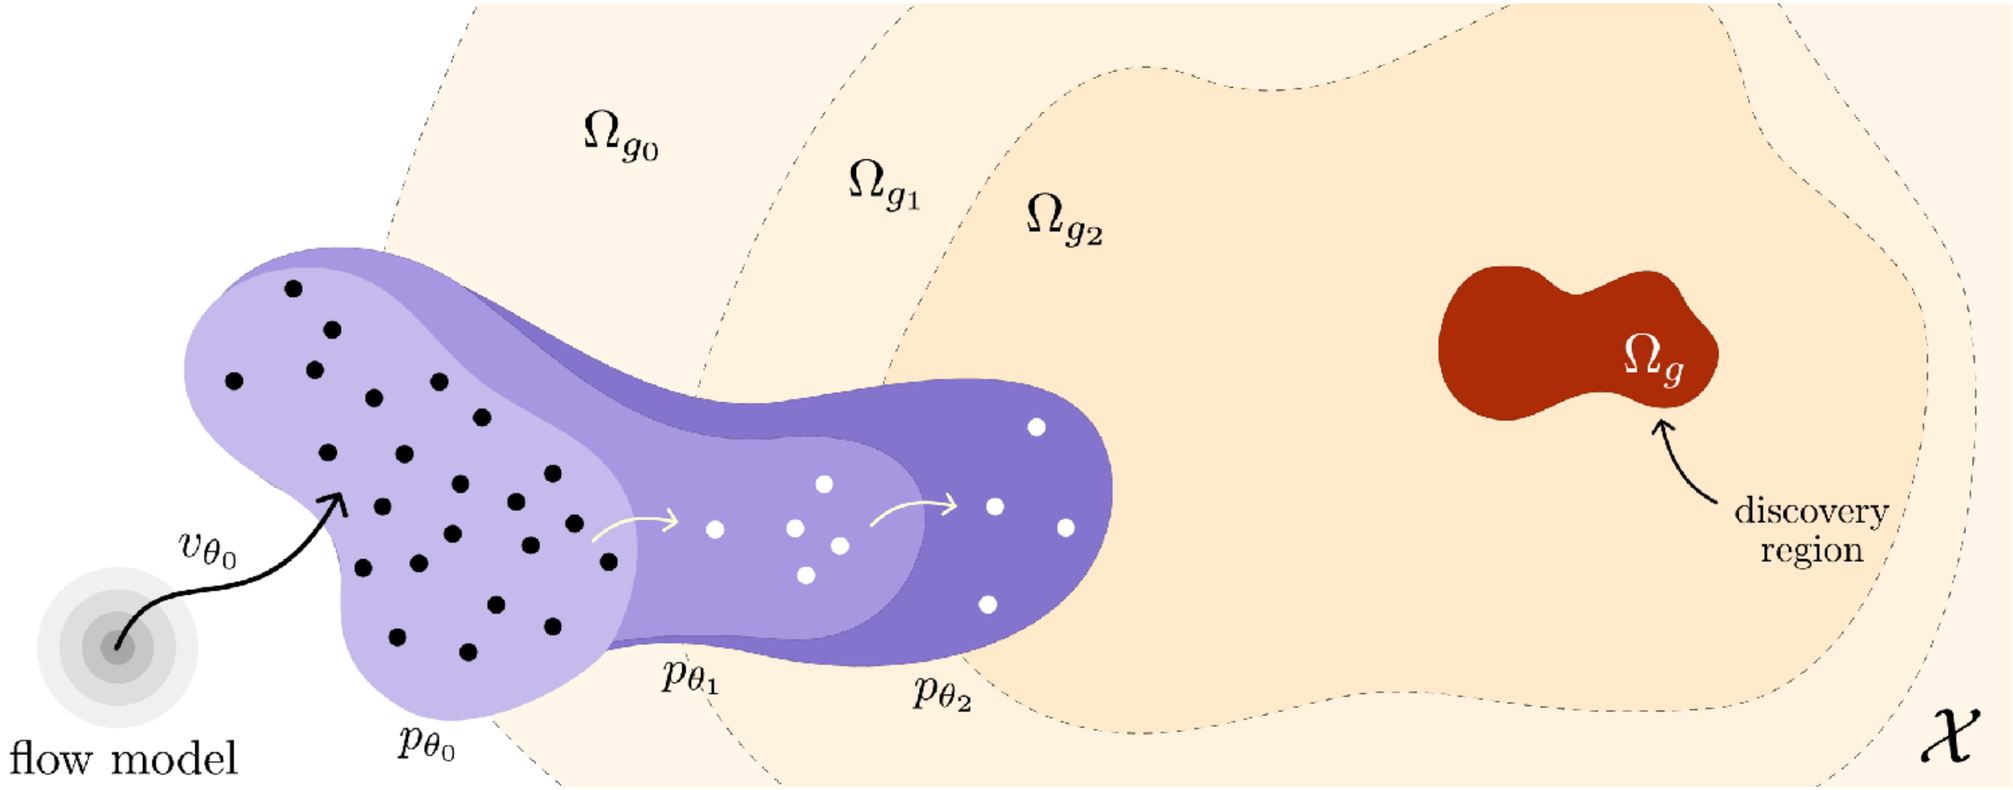

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

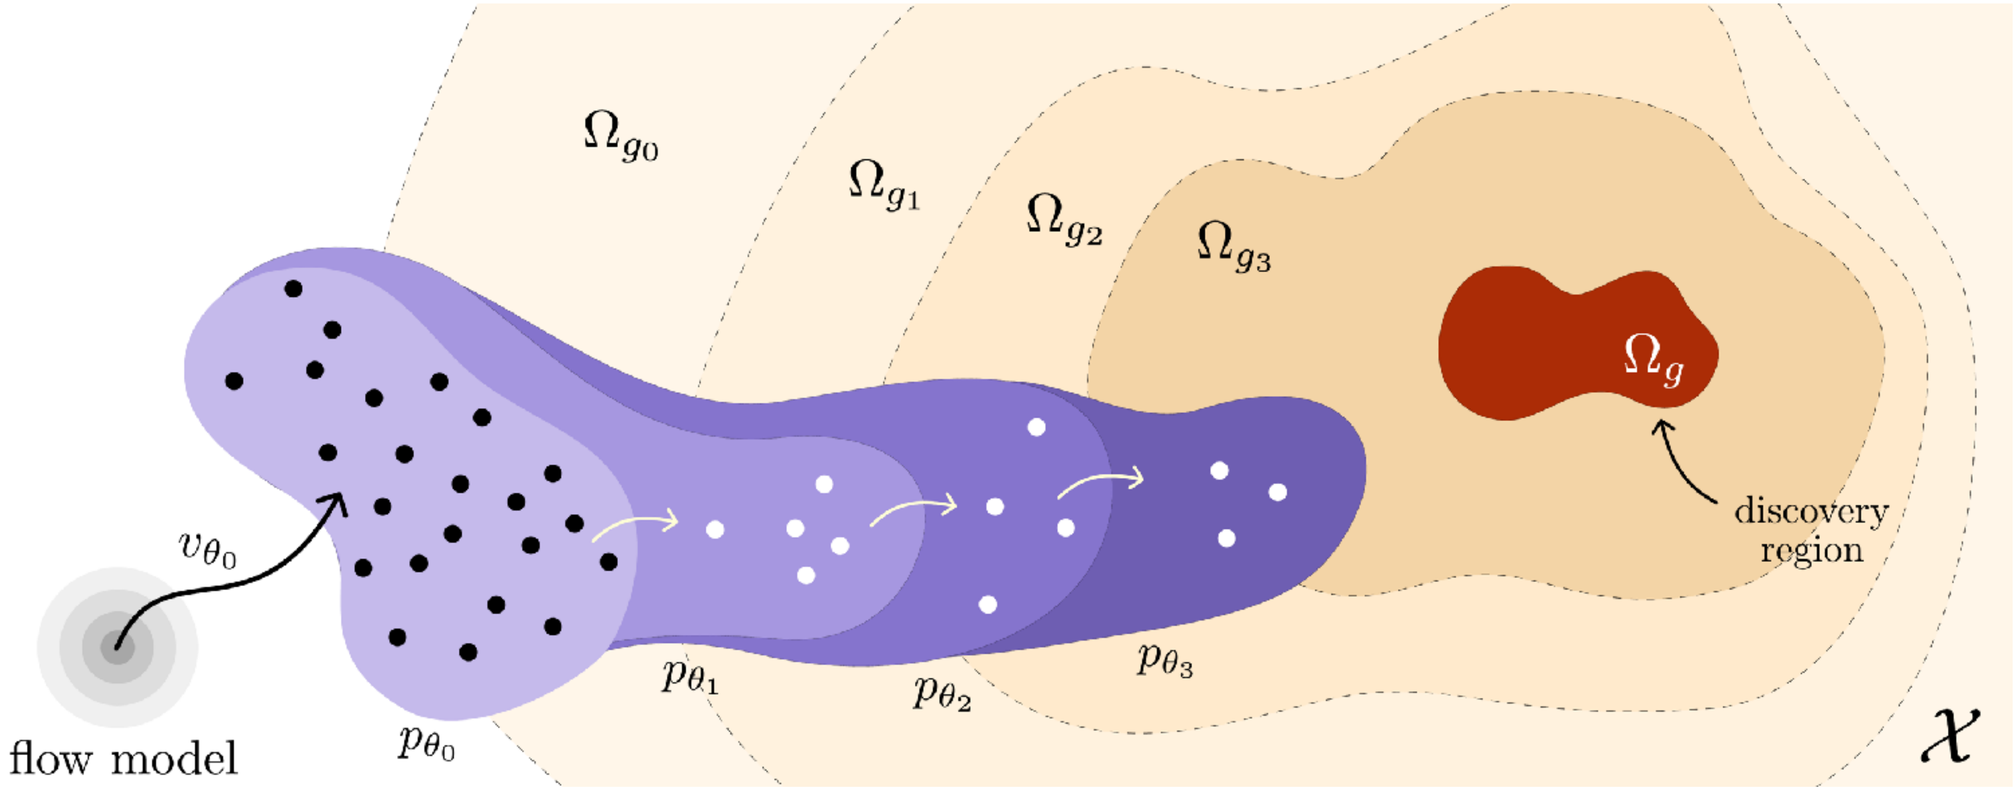

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

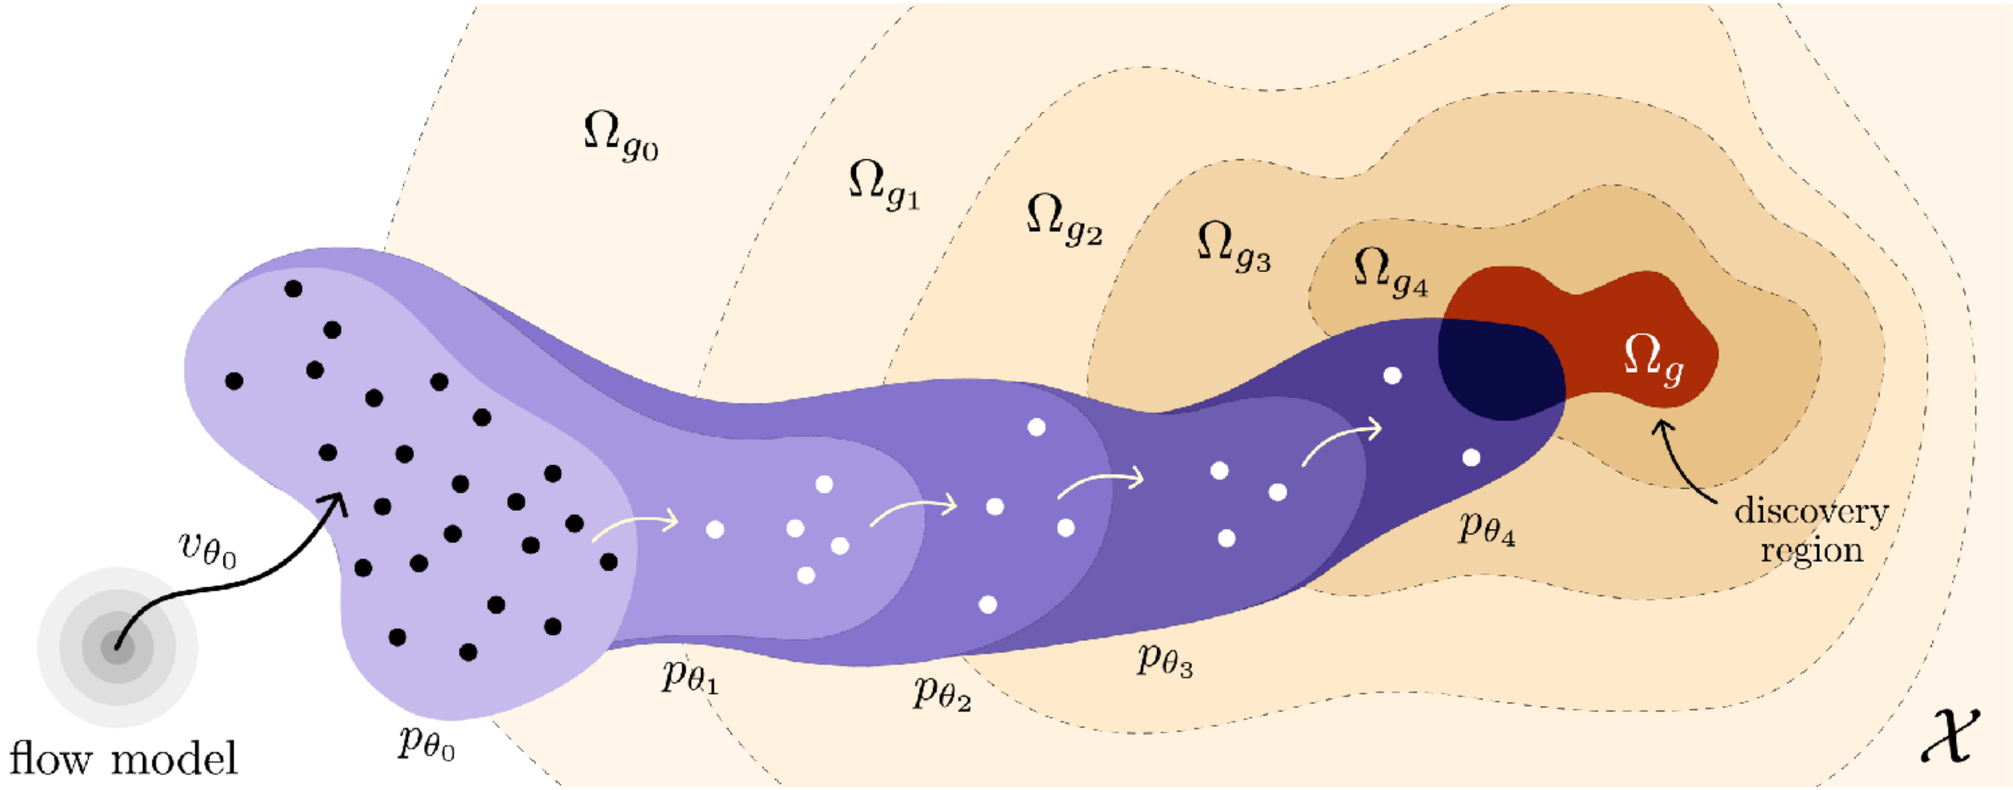

LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

$$\Omega_g := \{x \in \mathcal{X} : f_g(x) = 1\}$$

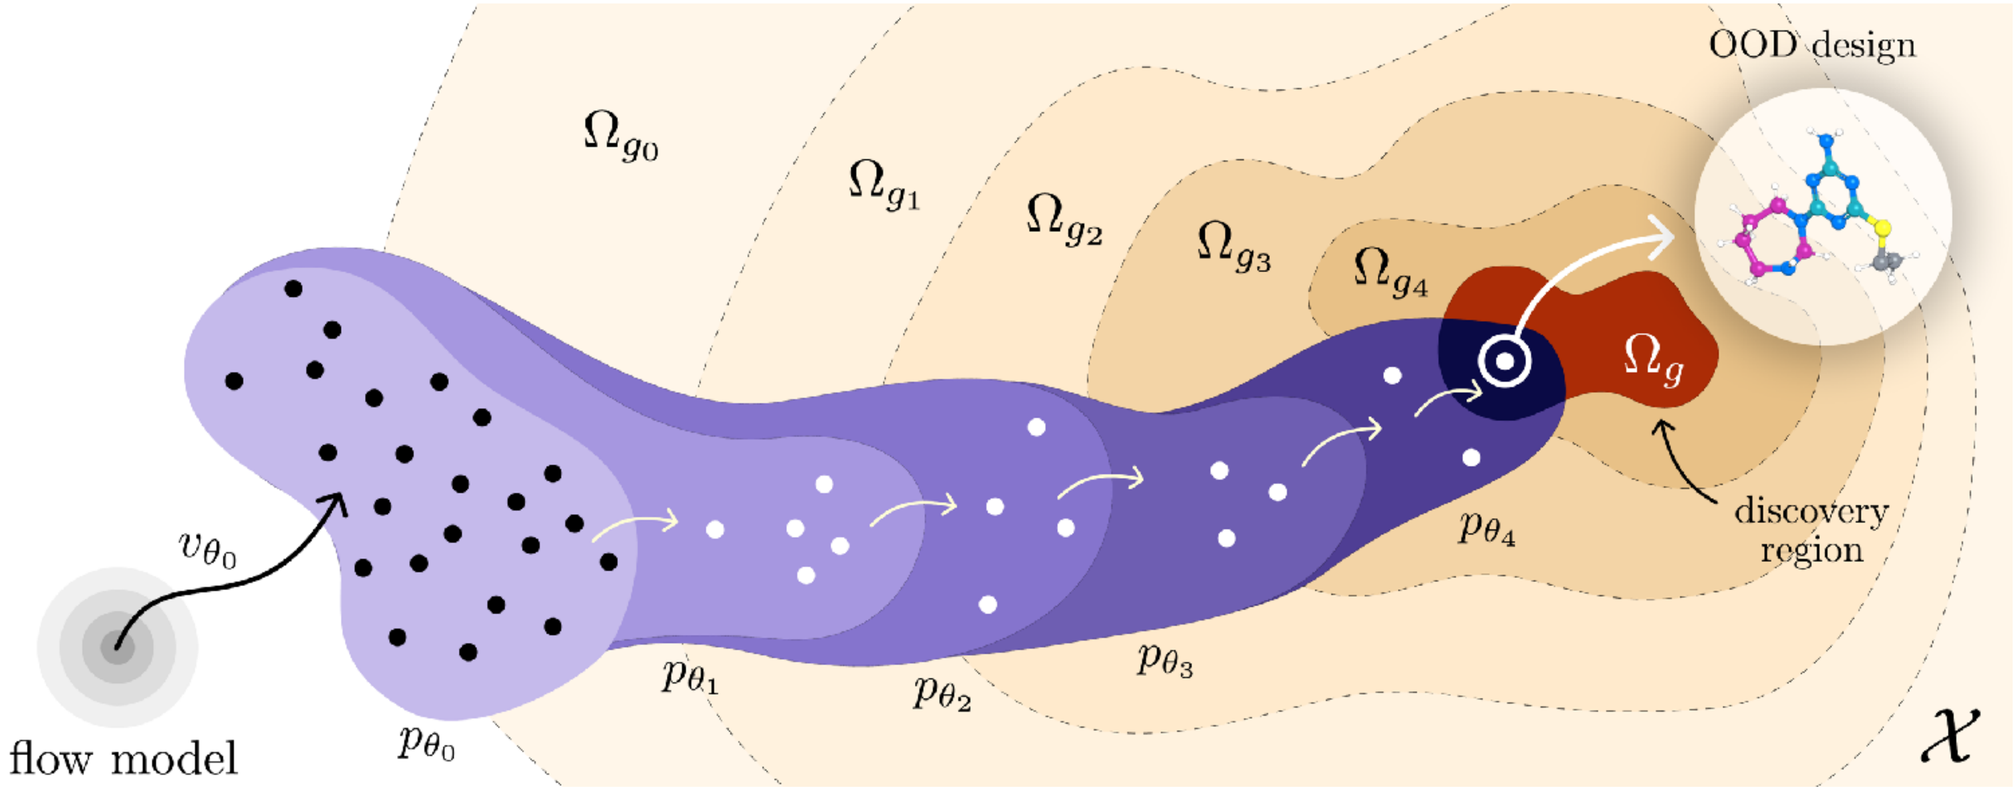

GOAL 1: Generate peptides with promising **anticancer activity** while maintaining **low toxicity** toward healthy cells.

LFD does not observe any anticancer proxy predictor!

GOAL 1: Generate peptides with promising **anticancer activity** while maintaining **low toxicity** toward healthy cells.

LFD does not observe any anticancer proxy predictor!

Method	Therapeutic Peptides (Anti-Cancer)		
	AntiCP2-alt. (%) ↑	HemoPI2 (%) ↓	Distance from prior ↑
Pre-trained	26.43 ± 2.91	**40.85** ± 1.89	0.00 ± 0.00
DPO	25.29 ± 3.19	42.77 ± 1.81	0.08 ± 0.04
DIFFUSIONNFT	33.26 ± 2.76	45.61 ± 2.39	0.16 ± 0.03
GUIDANCE	35.94 ± 2.33	47.25 ± 1.81	0.29 ± 0.03
LFD	**53.54** ± 3.32	42.48 ± 2.04	**0.31** ± 0.03

LFD surpasses state-of-the-art reward-guided methods on held-out evaluators.

53.5% (LFD) vs 35.9% (best reward-guided baseline) on a held-out anticancer predictor.

GOAL 1: Generate peptides with promising **anticancer activity** while maintaining **low toxicity** toward healthy cells.

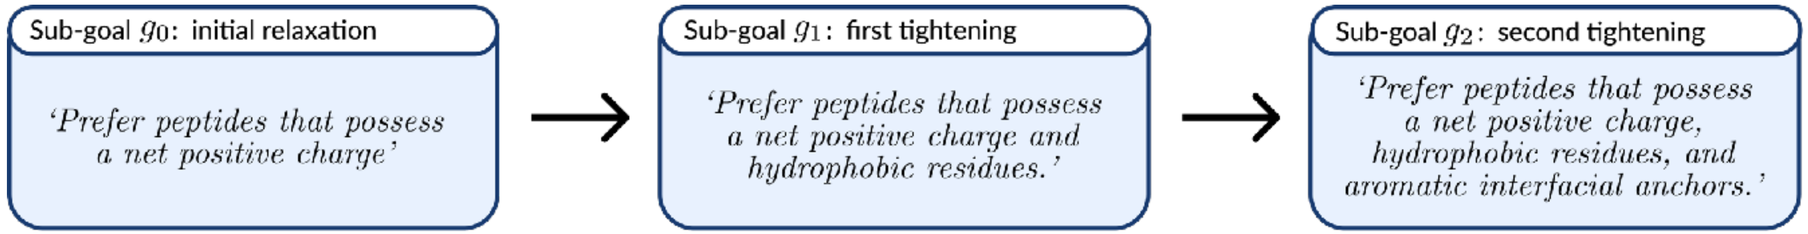

LFD generates sub-goals interpretable by scientists.

GOAL 1: Generate peptides with promising **anticancer activity** while maintaining **low toxicity** toward healthy cells.

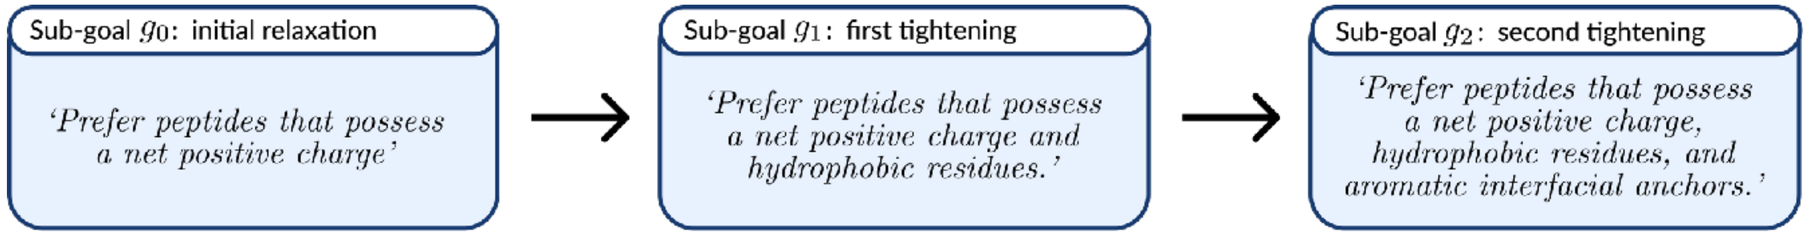

LFD generates sub-goals interpretable by scientists.

LLM interpretable traces

Peptide (−/−):
MLKGRWPDEQKKNHLMNYNFMLIVFVLTFNISFQSSNI
Charge index = +2.1; hydrophobicity = 47.4%; lower-third non-hemolysis band.

LLM rationale: *Violates the operative goal (g_3) by possessing a prolonged, contiguous hydrophobic tract (MLIVFVL) [...]*

LLM uncertainty: *In vitro hemolytic threshold and aggregation state in physiological buffer.*

INSIGHT: LFD generates scientist-interpretable sub-goals and decision-making traces, allowing human scientists to understand and steer the discovery process.

GOAL 3: Generate promising drug-like small molecules for **binding the dopamine D2 receptor**, balancing plausible receptor recognition with moderate lipophilicity and **chemically viable structures**.

LFD does not observe any docking proxy!

GOAL 3: Generate promising drug-like small molecules for **binding the dopamine D2 receptor**, balancing plausible receptor recognition with moderate lipophilicity and **chemically viable structures**.

LFD does not observe any docking proxy!

Method	Small Molecules (D2 Receptor)		
	GNINA ↑	Vina ↓	rDock ↓
Pre-trained	6.11 ± 0.33	−7.52 ± 0.39	−7.64 ± 1.33
DPO	5.60 ± 0.05	−6.81 ± 0.08	−8.11 ± 0.36
FLOW-GRPO	6.32 ± 0.06	−7.83 ± 0.09	−8.15 ± 0.27
DIFFUSIONNFT	6.23 ± 0.07	−7.90 ± 0.23	−9.39 ± 0.26
GUIDANCE	6.29 ± 0.07	−7.65 ± 0.21	−7.92 ± 0.25
LFD	**6.54** ± 0.09	**−8.48** ± 0.16	**−9.89** ± 0.38

LFD outperforms state-of-the-art reward-guided baselines on diverse **held-out docking scores.**

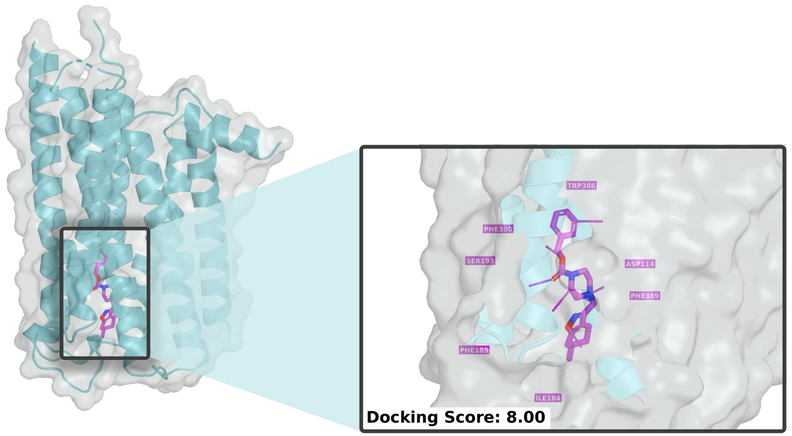

LFD-generated molecule docks to D2 dopamine receptor

98% **Validity** 89.5% **Conformer Quality**

GOAL 5: Generate **neutral**, connected small molecules for a kinase-focused discovery campaign that contain a **six-membered aza-aromatic ring** bearing a **free primary amino group** on a **ring carbon adjacent to a pyridine-like ring nitrogen**, and a distinct **fully saturated ring** containing a **free secondary amine**.

GOAL 5: Generate **neutral**, connected small molecules for a kinase-focused discovery campaign that contain a **six-membered aza-aromatic ring** bearing a **free primary amino group** on a **ring carbon adjacent to a pyridine-like ring nitrogen**, and a distinct **fully saturated ring** containing a **free secondary amine**.

Pretrained FlowMol3 **never** generated a molecule with all the required motifs in **38.500+ samples**.

↓

Desired structure is **out of** prior model's **generable set**

GOAL 5: Generate **neutral**, connected small molecules for a kinase-focused discovery campaign that contain a **six-membered aza-aromatic ring** bearing a **free primary amino group** on a **ring carbon adjacent to a pyridine-like ring nitrogen**, and a distinct **fully saturated ring** containing a **free secondary amine**.

Pretrained FlowMol3 **never** generated a molecule with all the required motifs in **38.500+ samples**.

→

LFD finds a hit in a 512-sample batch at round 6.

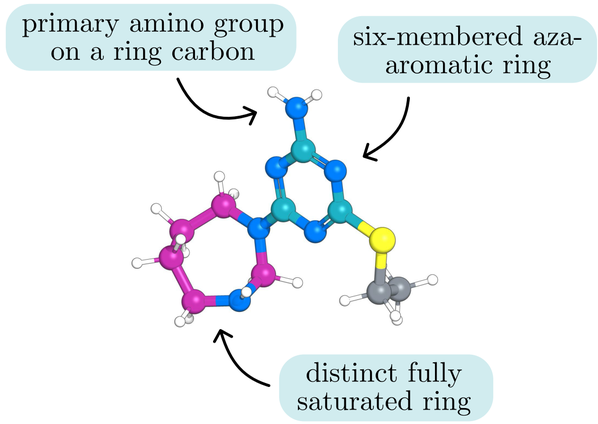

Generative Discovery

Learning to generate **diverse and rare** designs, **expanding beyond the prior**, toward **high-level discovery goals**.

Distributional Flow Control for Discovery

Which candidates matter?

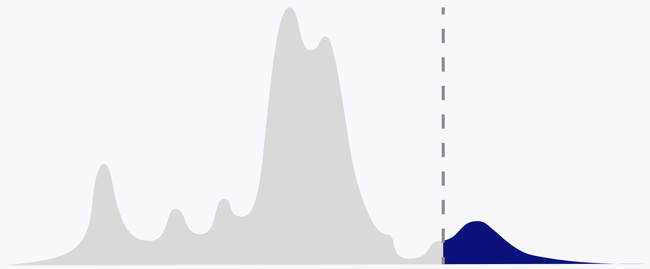

→ FDC

Out-of-Distribution Flow Modeling

Which regions can we reach?

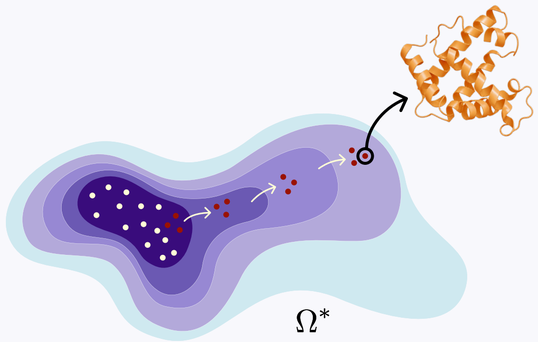

→ ActFlow

Sample-Efficient Discovery via LLM Distillation

What feedback guides us there?

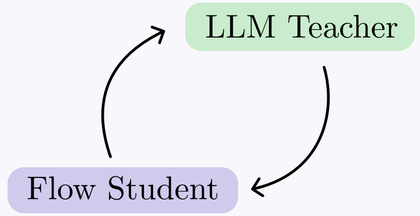

→ LFD

Mathematical Foundations of Exploration

Generative Discovery Algorithms

Experimental Sciences (Wet-lab Discovery)

Enzymes







Caltech

Therapeutic Peptides

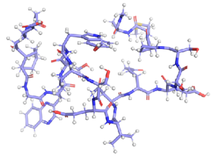



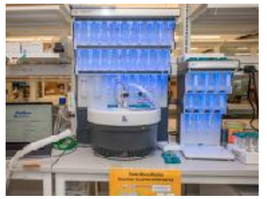

Penn
University of Pennsylvania

Functional Materials

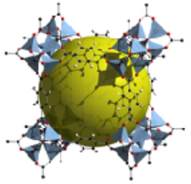





ETH zürich

Thanks!

Collaborators

Kimon Protopapas, Ya-Ping Hsieh, Sven Gutjahr, Luca Schaufelberger, Zifan Wang, Xiaoyu Mo, Malte Franke, Cristian Perez Jensen, Andrey Kharitenko, Zebang Shen, Marin Vlastelica, Guy Schacht, Mojmir Mutny, Ziyad Sheebaelhamd, Bruce Lee, Sophia Tang, Cheng-Hao Liu, Federico Di Gennaro, Angelo Gnazzo, Michael M. Zavlanos, Harini Vijayshankar, Dylan Abramson, Basile Wicky, Karl H. Johansson, Fanny Yang, İsmail İlkan Ceylan, Kjell Jorner, Florian Dörfler, Frances H. Arnold, Pranam Chatterjee, Michael M. Bronstein, Yisong Yue, Niao He, Andreas Krause.

ETH zürich ETH AI CENTER

Caltech

Thanks!

Papers + Code →

riccardodesanti.com/research

Connect:

X: @r_de_santi

email: rdesanti@ethz.ch

Building the foundations of **discovery systems** that learn to **explore beyond the data**.

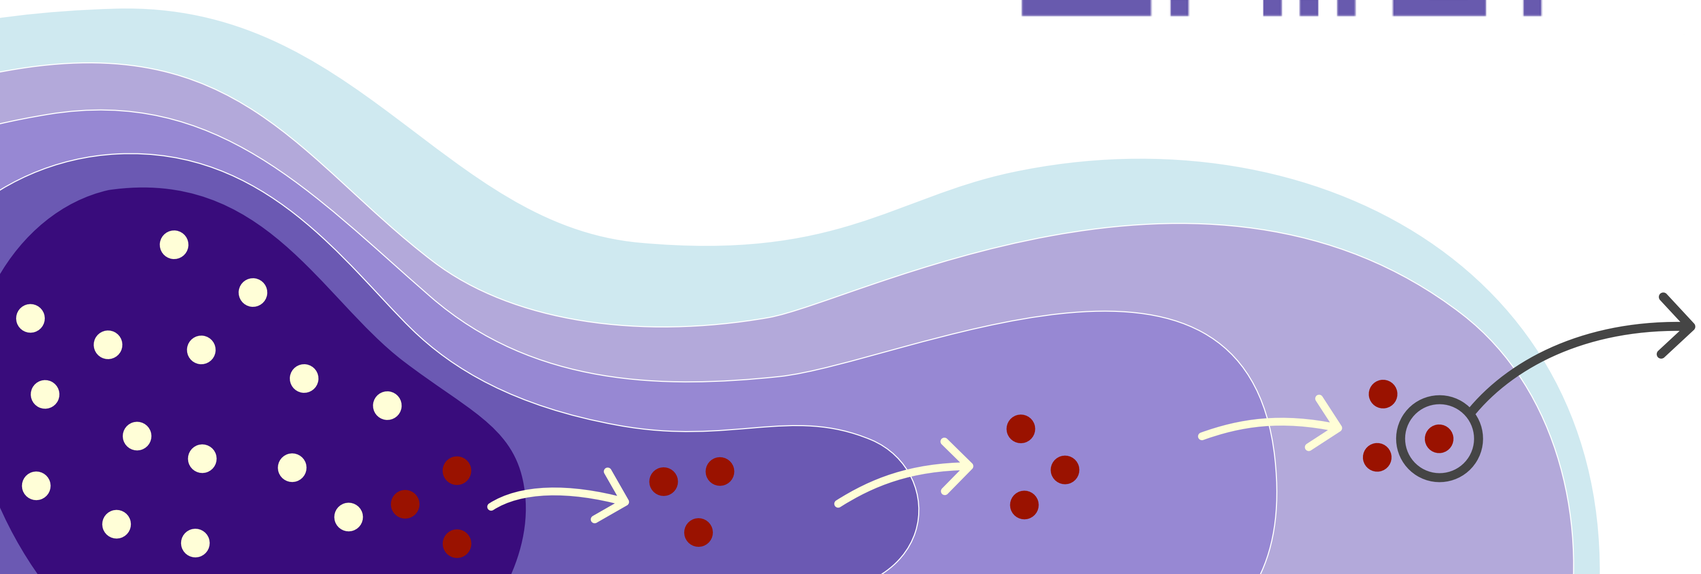



Riccardo De Santi