ETH zürich ETH AI CENTER

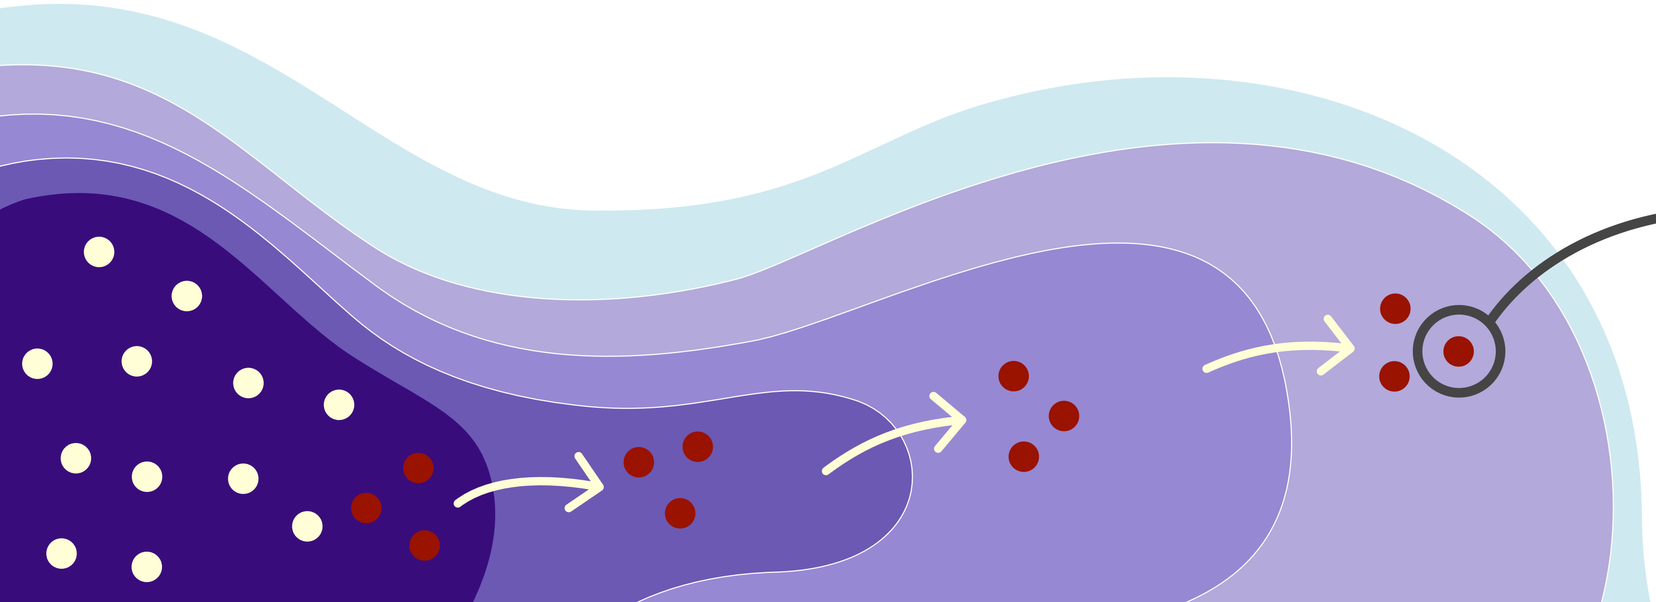
Caltech

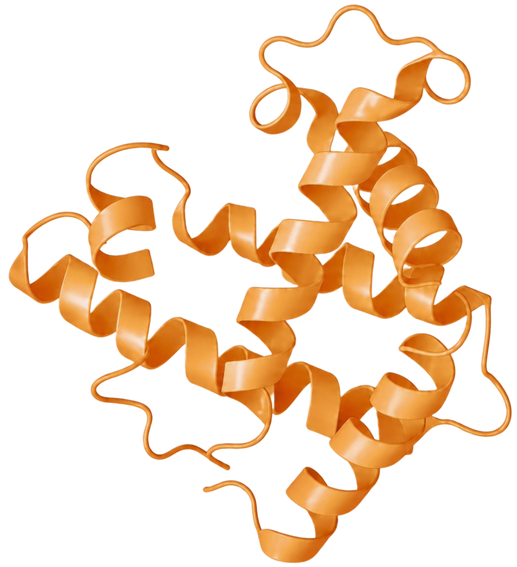
Generative Discovery Beyond the Data

Riccardo De Santi

PhD Candidate at ETH Zurich

@r_de_santi rdesanti@ethz.ch





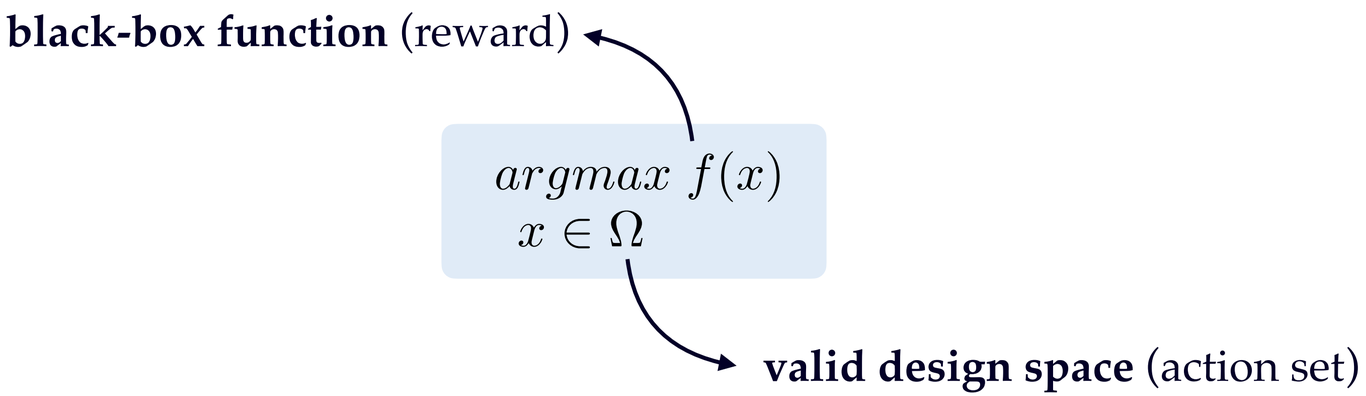
Rising Stars in DS, November 10th, 2026

What is a Discovery Problem?

A real-world discovery task:

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

What is a Discovery Problem?

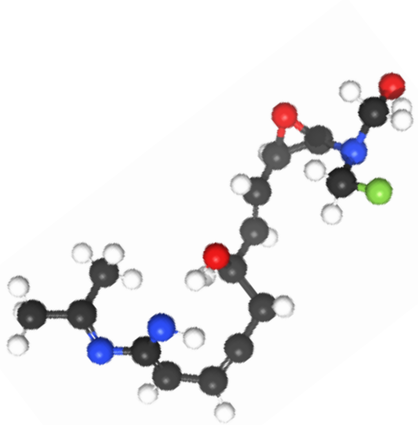
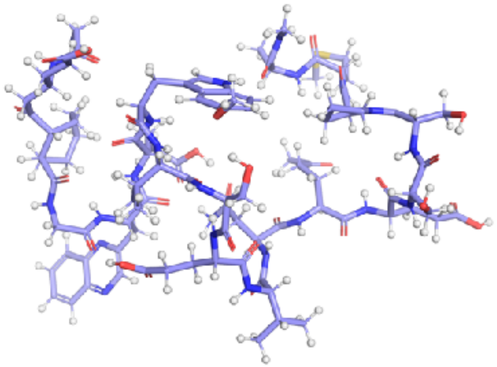
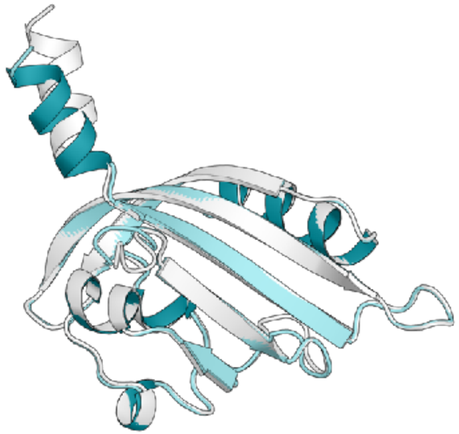
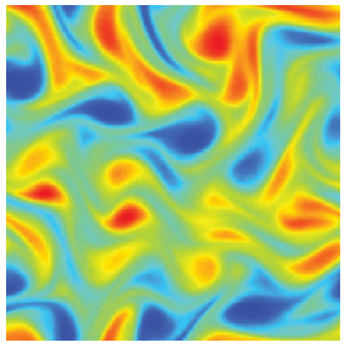
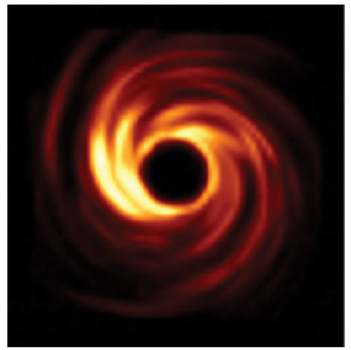


What is a Discovery Problem?

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)

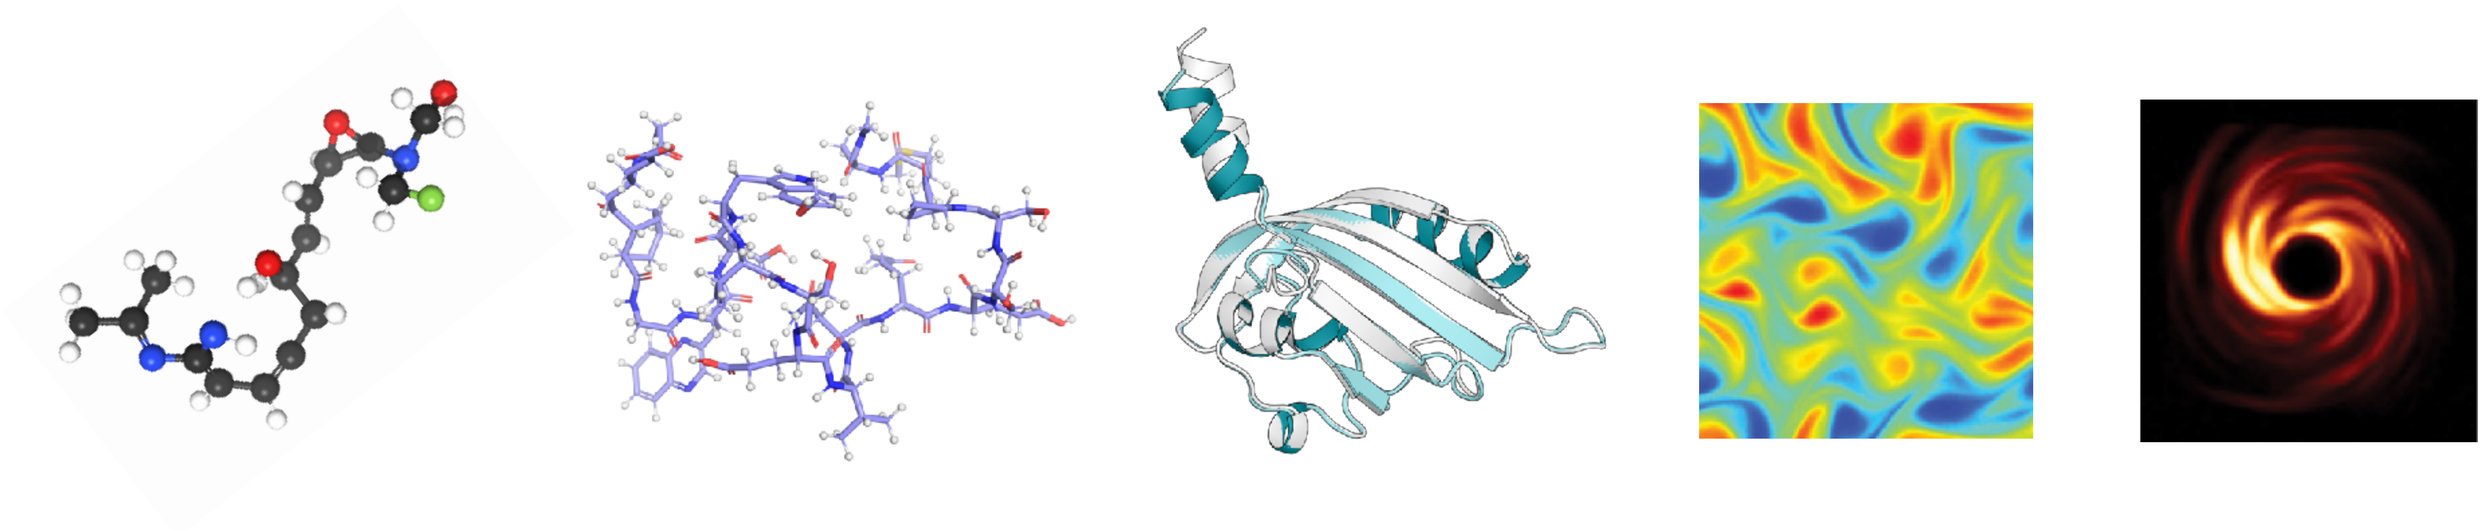
$$argmax\ f(x)$$
$$x \in \Omega$$

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)

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

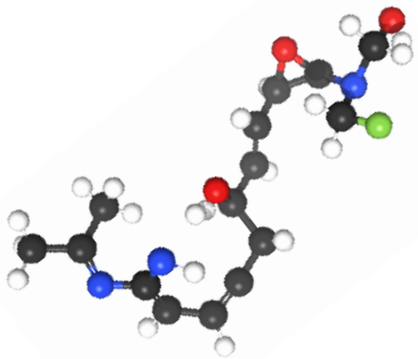
valid design space (action set)

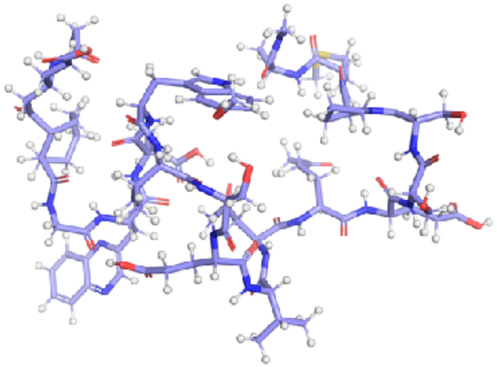
$\approx 10^{60}$ drug discovery

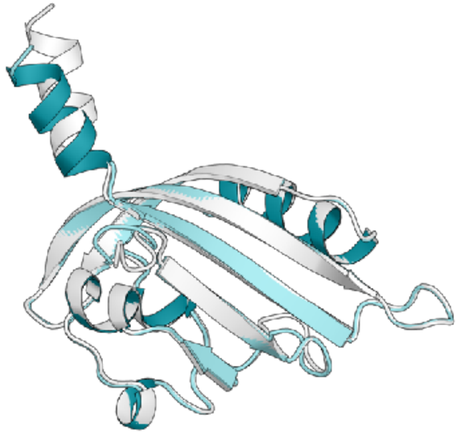
$\approx 10^{80}$ material design

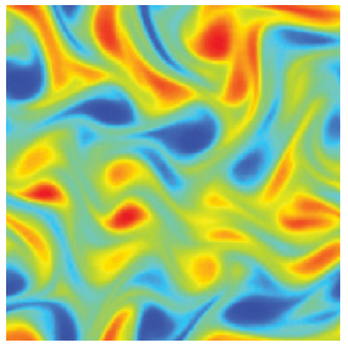
f and Ω are **unknown** and **hard to express**

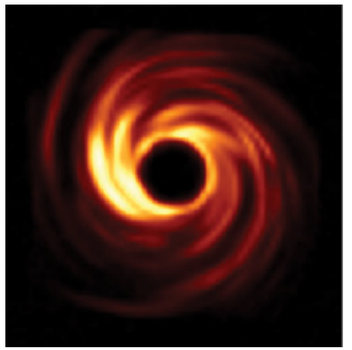
→ Generative Discovery







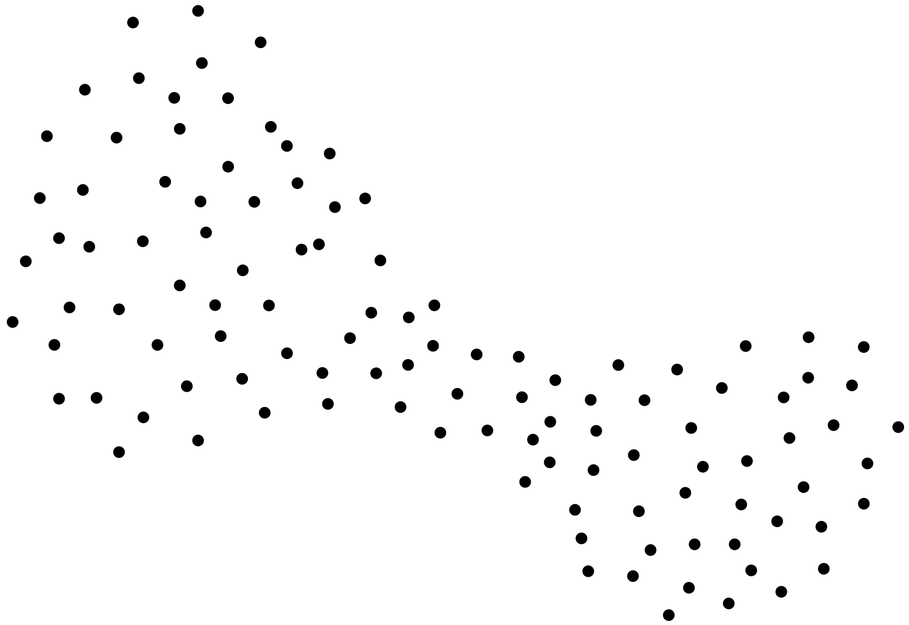




[Zheng et al., 2025]

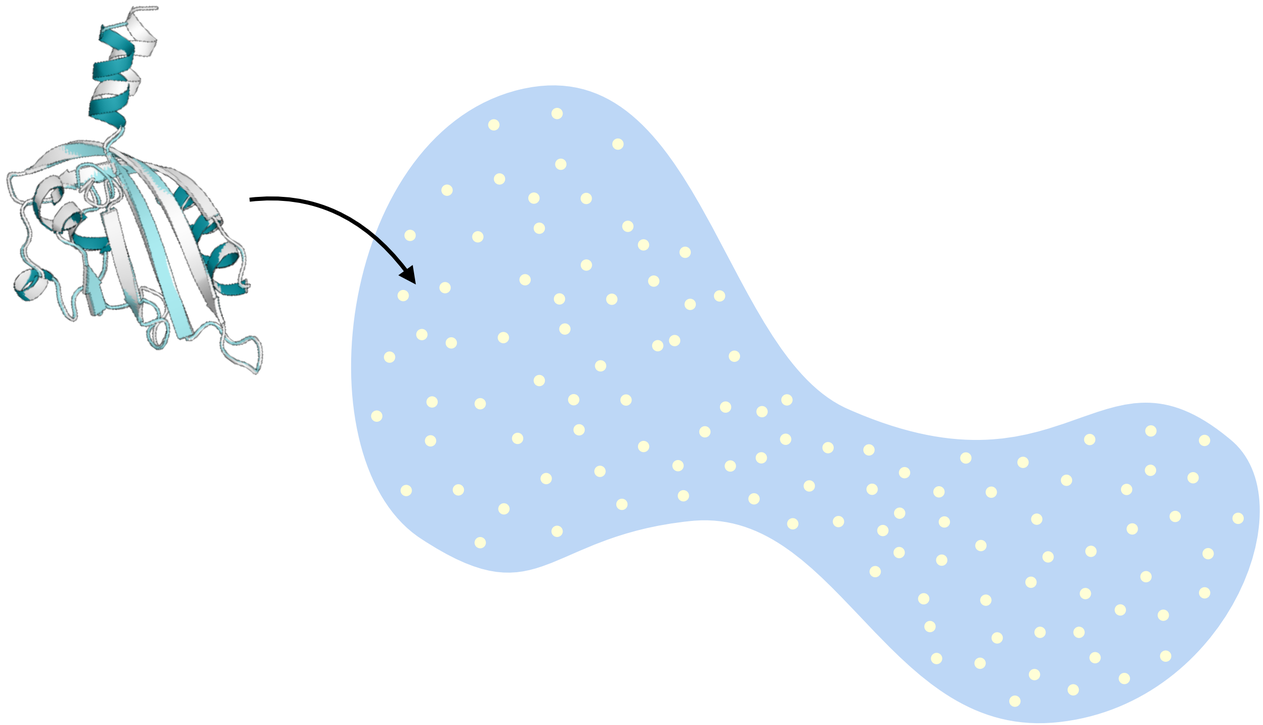
Generative Discovery Paradigm

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



Generative Discovery Paradigm

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

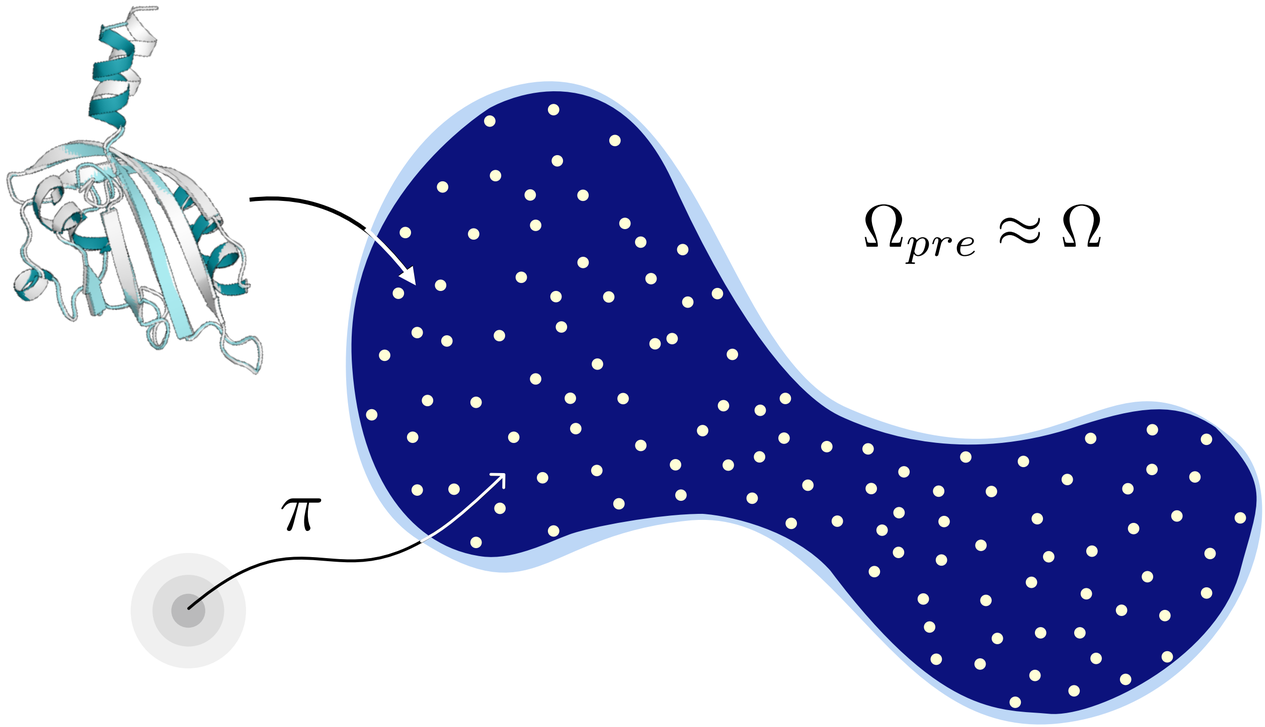


Data $D = \{x_i\}_{i=1}^{n}$

Generative Discovery Paradigm

valid design space Ω

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



$\Omega_{pre} \approx \Omega$

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

↓

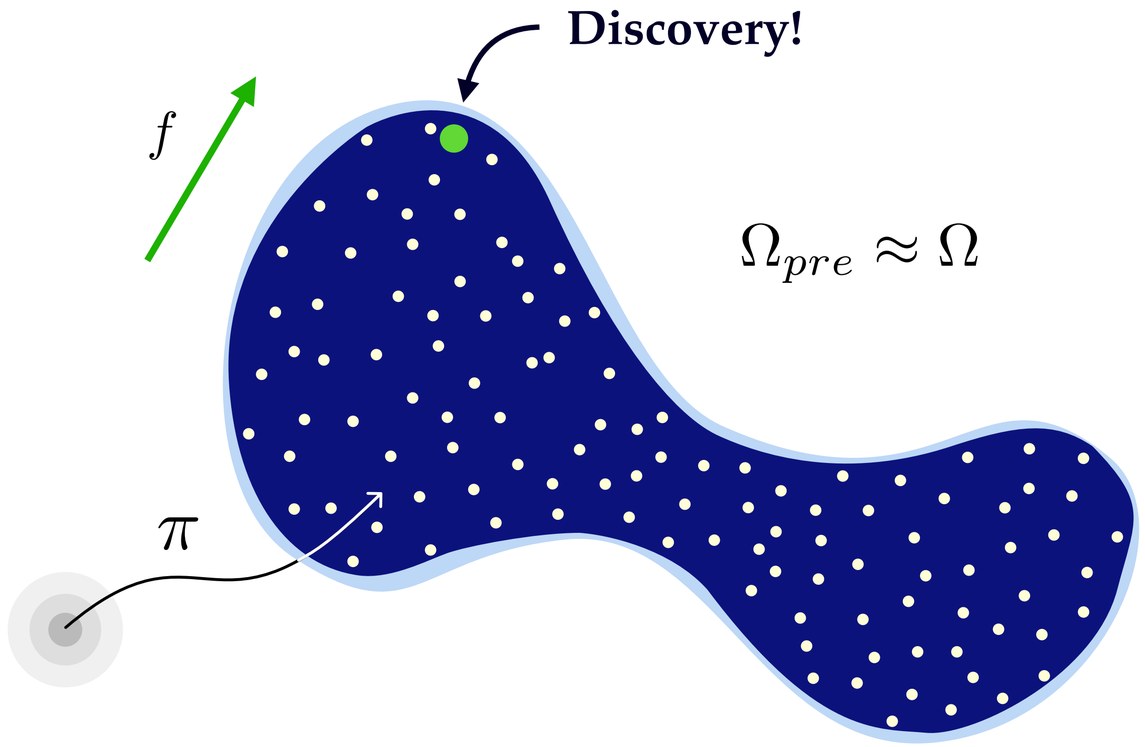
Pre-trained generative model π

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}



$\Omega_{pre} \approx \Omega$

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

↓

Pre-trained generative model π

↓

Adapt π to optimize f

(e.g., via Guidance, RL)

Discovery problem

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

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

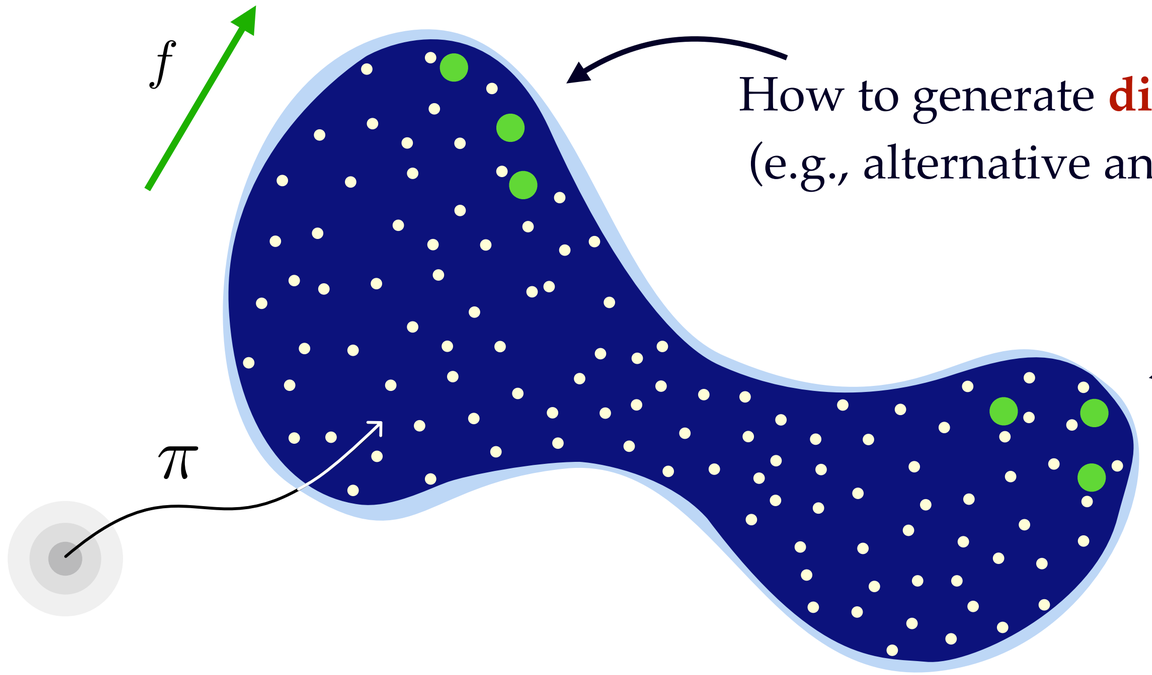
Fundamental Challenge of Generative Discovery

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)

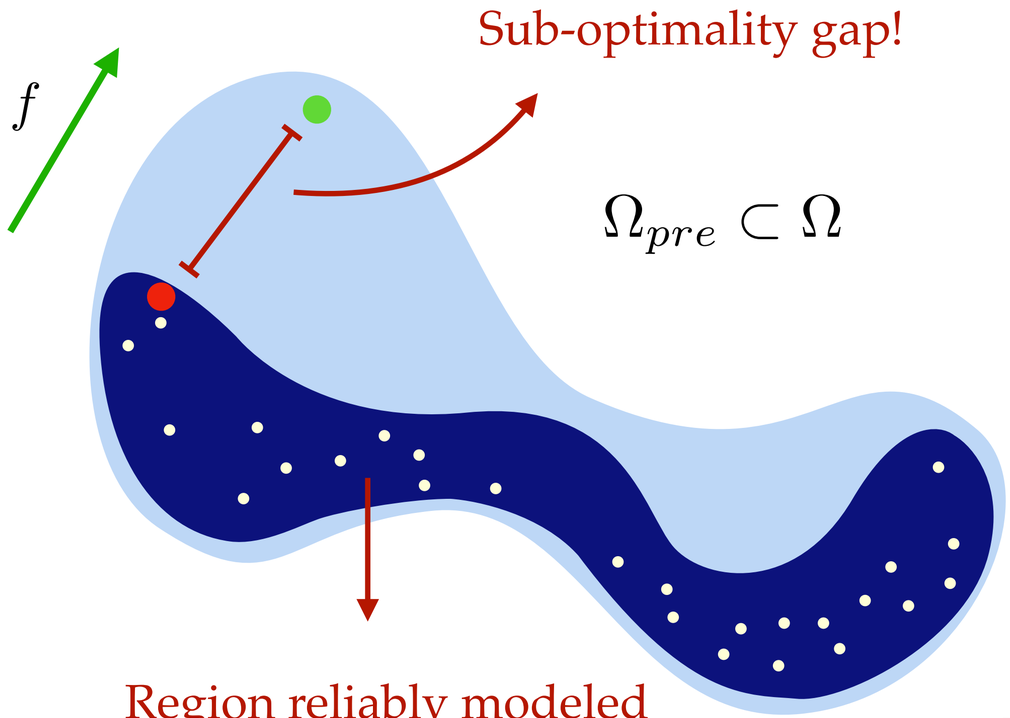
Fundamental Challenge of Generative Discovery

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

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

classic optimization
(e.g., convex optimization)

f known

Ω known

Fundamental Challenge of Generative Discovery

$$\begin{aligned} &argmax\ f(x) \\ &x \in \Omega \end{aligned}$$

classic optimization (e.g., convex optimization)		**decision-making** (e.g., bandits, BO, RL)	
f	known	f	learned
Ω	known	Ω	known

Fundamental Challenge of Generative Discovery

$$\begin{aligned} &argmax\ f(x) \\ &x \in \Omega \end{aligned}$$

classic optimization (e.g., convex optimization)		**decision-making** (e.g., bandits, BO, RL)		**generative discovery** (e.g., drug design)	
f	known	f	learned	f	learned
Ω	known	Ω	known	Ω	learned

Fundamental Challenge of Generative Discovery

$$\begin{aligned} &argmax\ f(x) \\ &x \in \Omega \end{aligned}$$

classic optimization (e.g., convex optimization)		**decision-making** (e.g., bandits, BO, RL)		**generative discovery** (e.g., drug design)	
f	known	f	learned	f	learned
Ω	known	Ω	known	Ω	learned

Takeaway

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

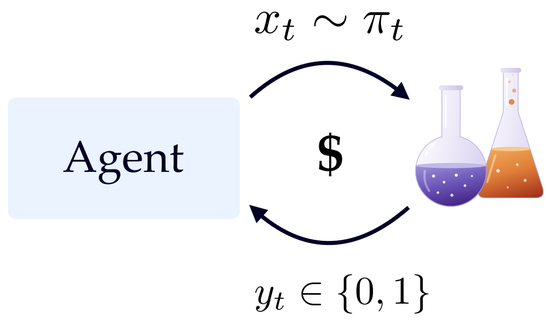
Fundamental Challenge of Generative Discovery

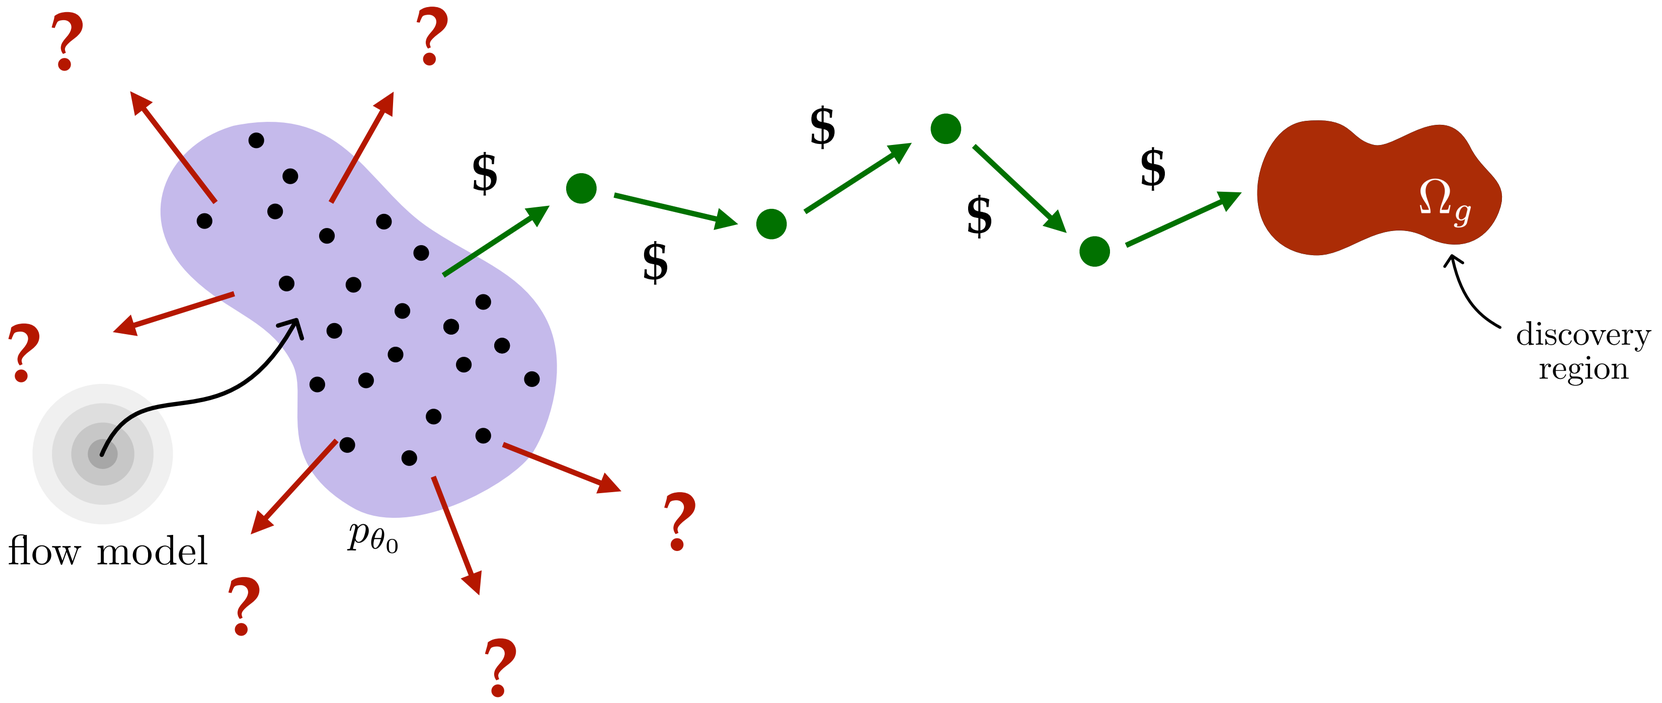
Discovery problem

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

High-level discovery goal

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





Generative Discovery

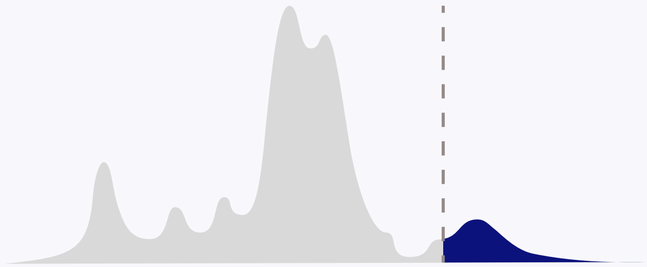
Learning to generate **diverse and rare** designs, **expanding beyond the prior**, toward **high-level discovery goals**.

Generative Discovery

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

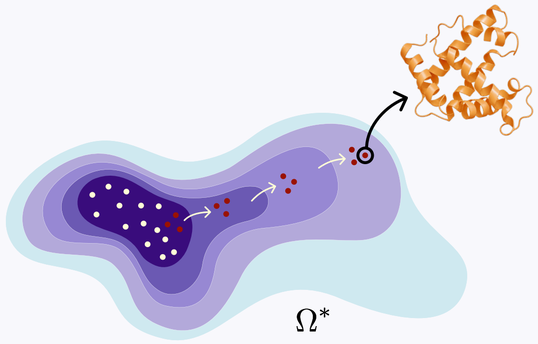
Distributional Flow Control for Discovery

Which candidates matter?



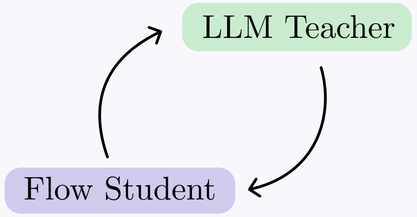
Out-of-Distribution Flow Modeling

Which regions can we reach?

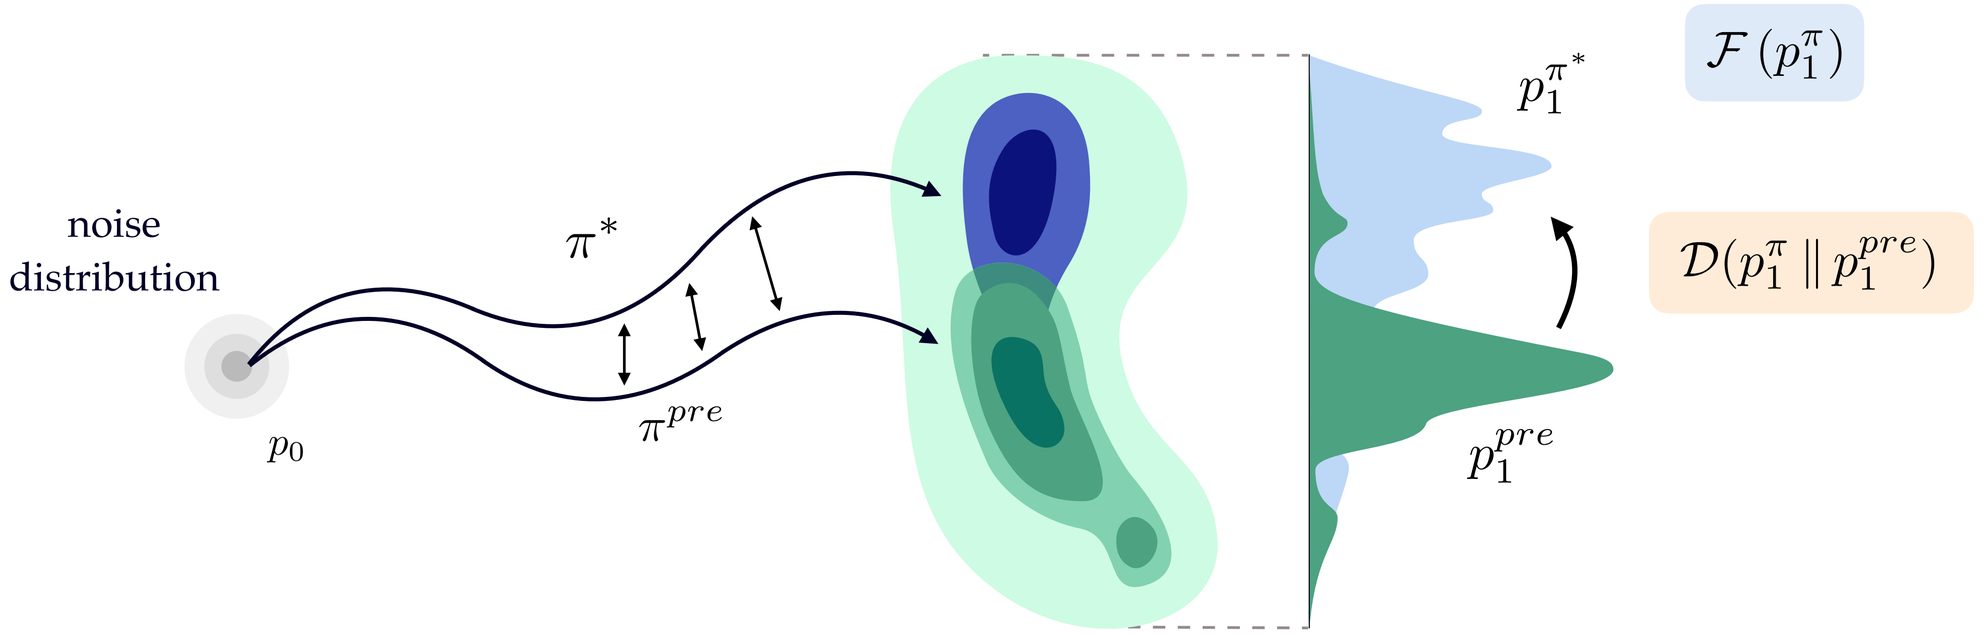


Sample-Efficient Discovery via LLM Distillation

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

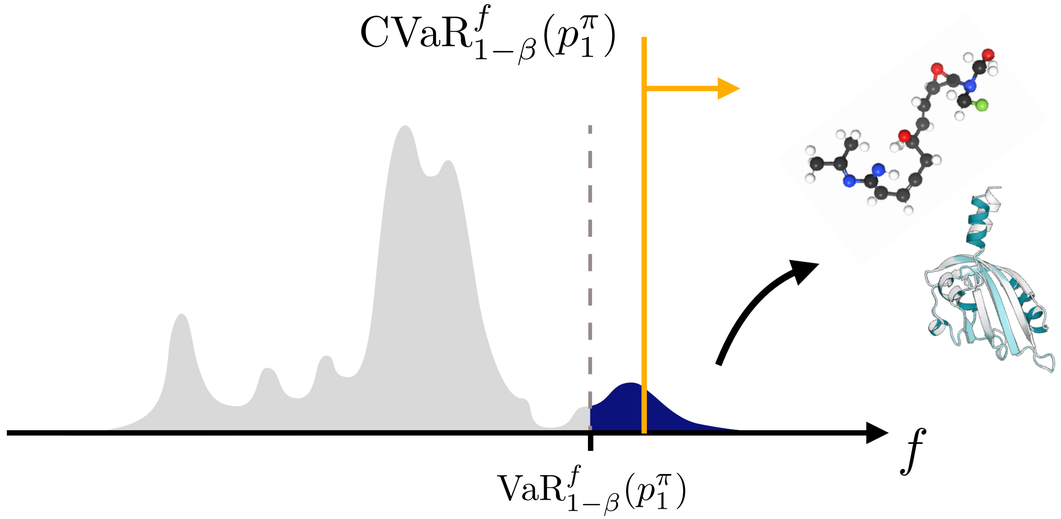
$$\pi^* \in \operatorname{argmax} \; \mathcal{F}(p_1^\pi) - \alpha \mathcal{D}(p_1^\pi \,\|\, p_1^{pre})$$

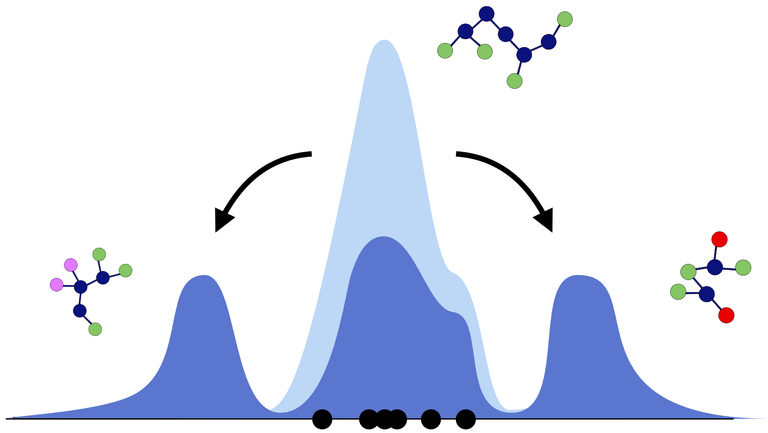
Extreme event generation

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

Diverse generation

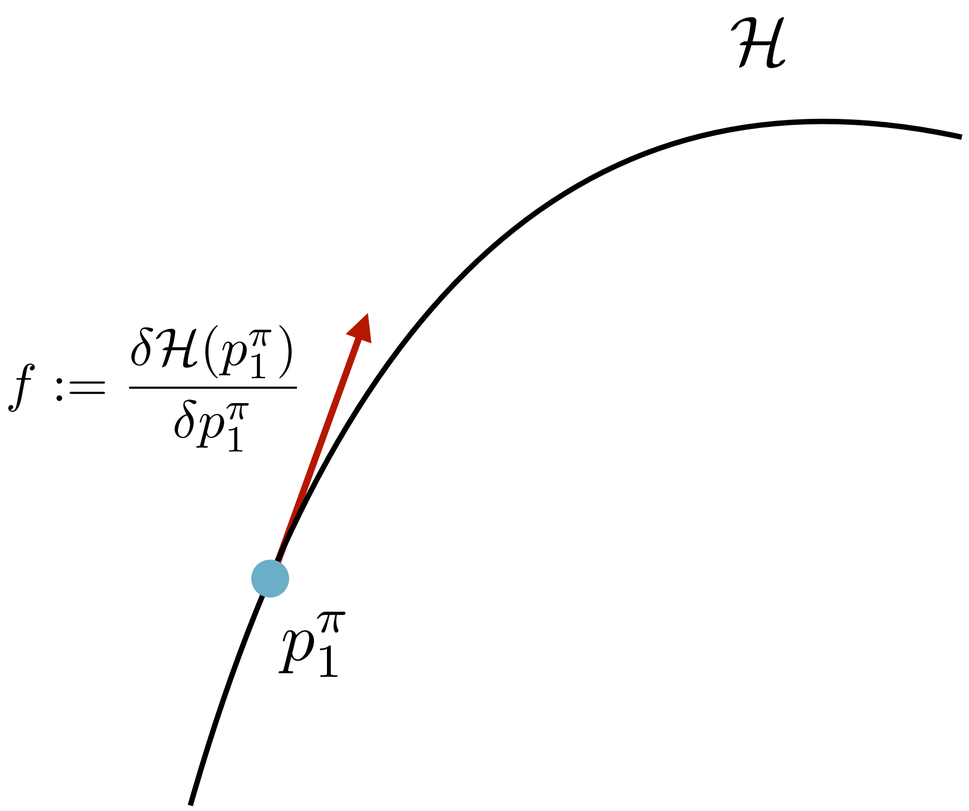
$$\mathcal{H}(p_1^\pi) := -\mathbb{E}_{x\sim p^\pi}[\log p^\pi(x)]$$





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

Flow Density Control (FDC)



Entropy First Variation (Surprise)

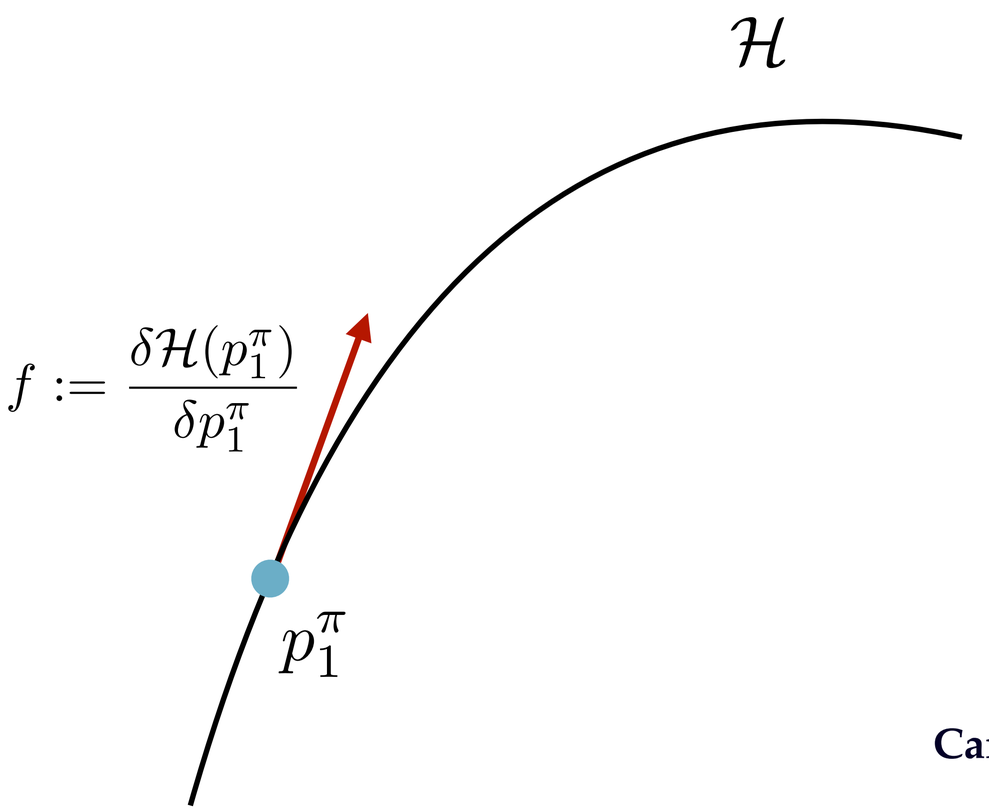
$$f := \frac{\delta \mathcal{H}\left(p_1^\pi\right)}{\delta p_1^\pi} = -\log\left(p_1^\pi\right)$$

surprise as **intrinsic reward function**

Log likelihoods are hard to estimate!

(currently does not scale)

Flow Density Control (FDC)



Entropy First Variation (Surprise)

$$f := \frac{\delta \mathcal{H}\left(p_1^\pi\right)}{\delta p_1^\pi} = -\log\left(p_1^\pi\right)$$

surprise as **intrinsic reward function**

Gradient of Surprise is the Negative Score

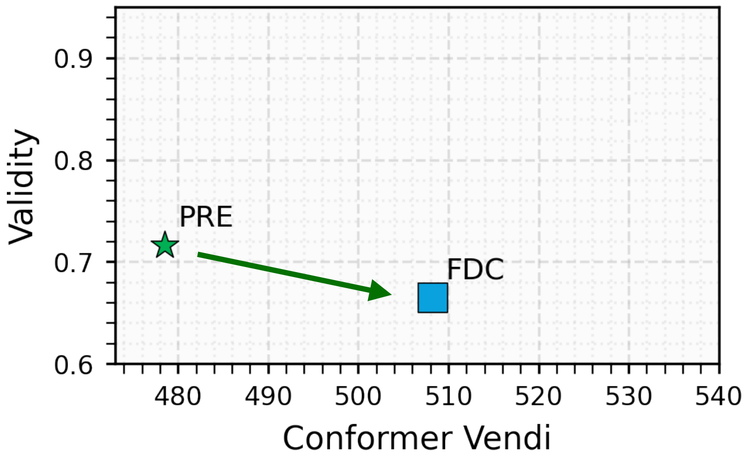
$$\nabla f = -\nabla \log\left(p_1^\pi\right) \approx -s^\pi(\cdot, 1)$$

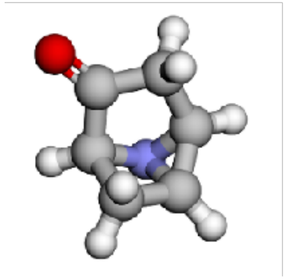
score neural approximator (e.g., *denoiser* network)

Can use any gradient-based FT method.

High-dimensional exploration with no estimation process!

Diverse Generation via Entropy

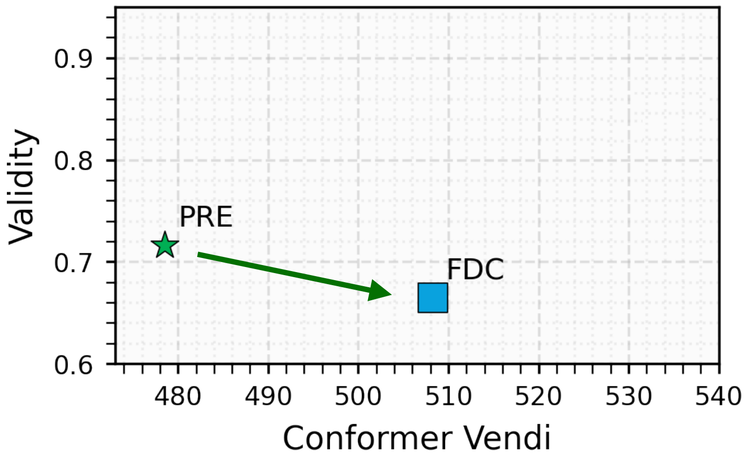


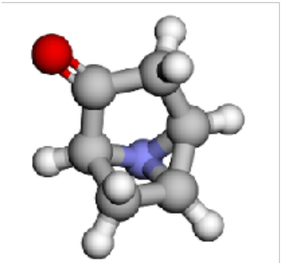


FDC increases structural diversity
while maintaining **high validity**

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

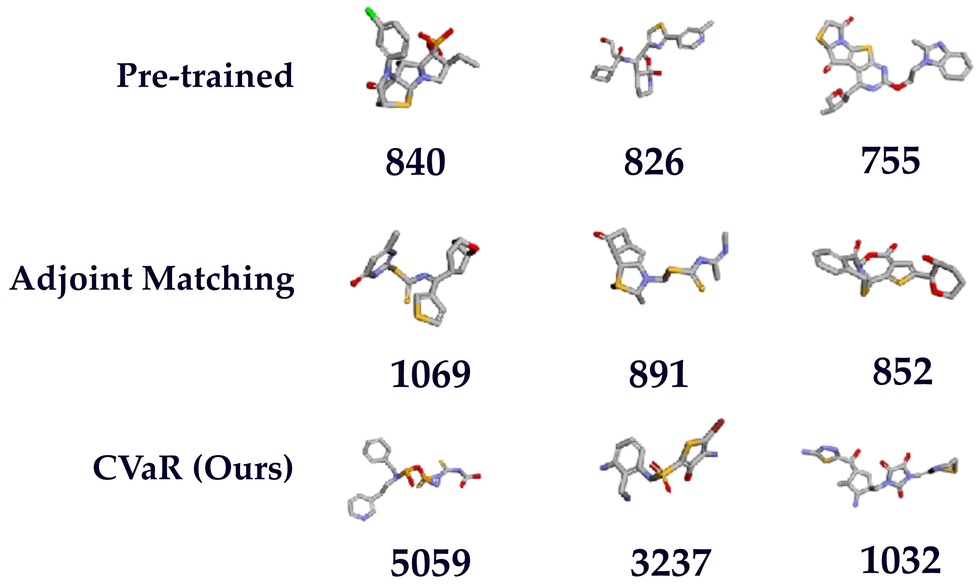
Diverse Generation via Entropy





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

Extreme Event Generation via CVaR



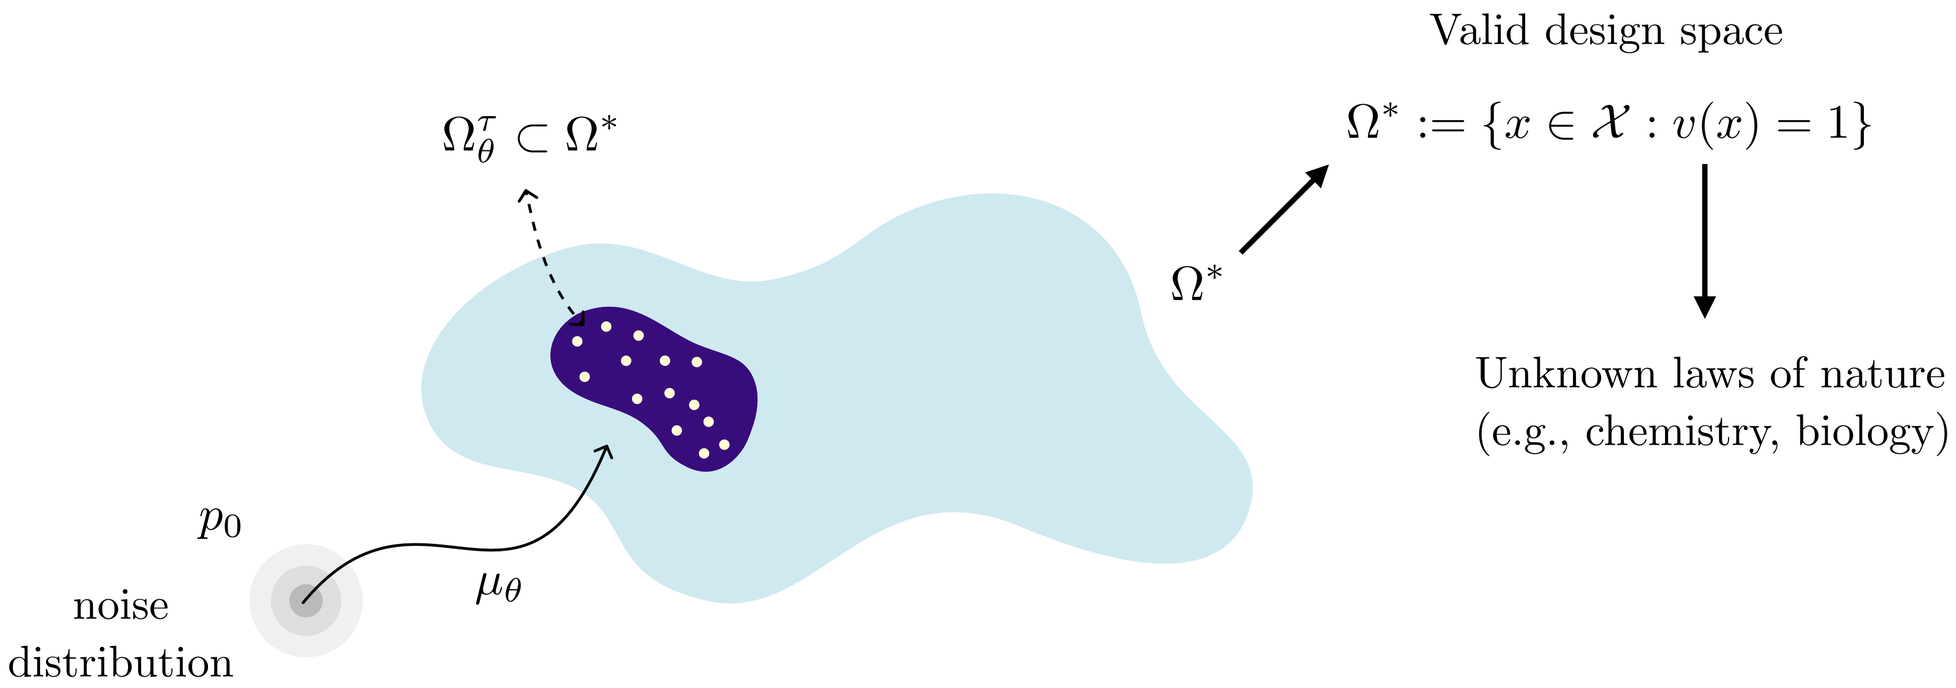
~5x higher top-sample reward

FDC significantly **improves top-samples**

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

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

Out-of-Distribution Generative Modeling

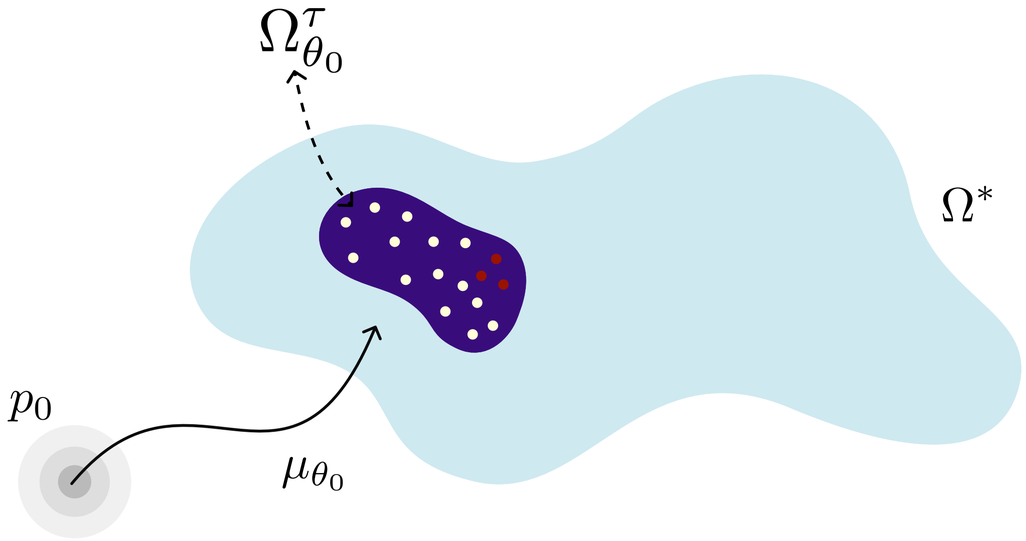


Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

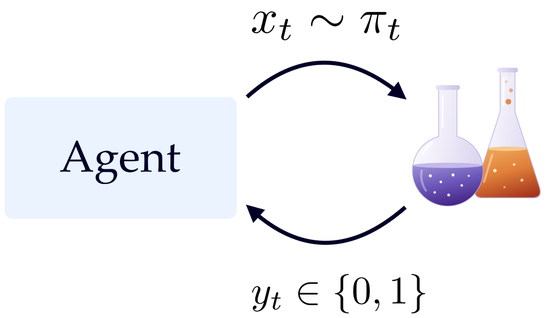
$$\Omega_\theta^\tau := \{x \in \mathcal{X} : p_1^\theta(x) \geq \tau\}$$

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

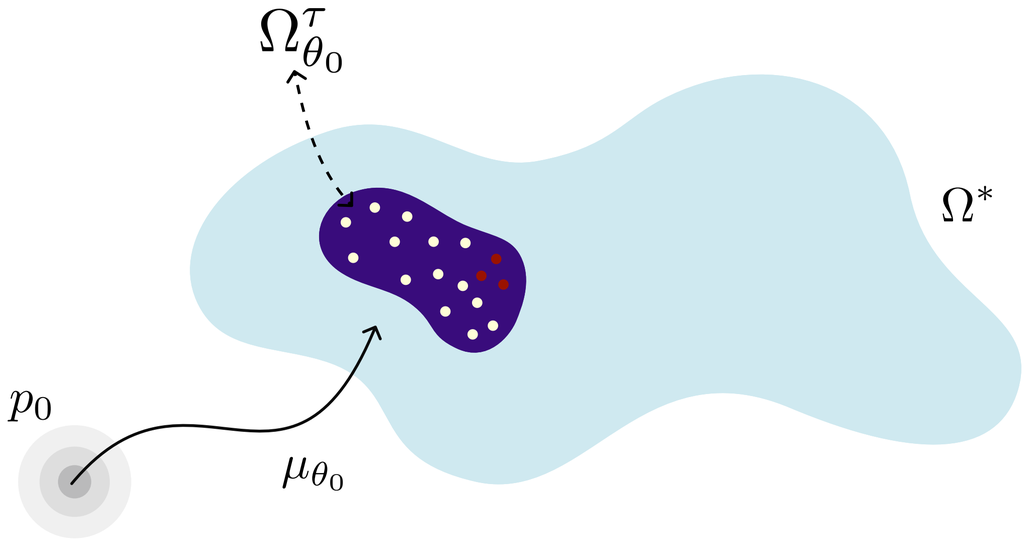
Generable Set Expansion: Beyond Distribution Matching



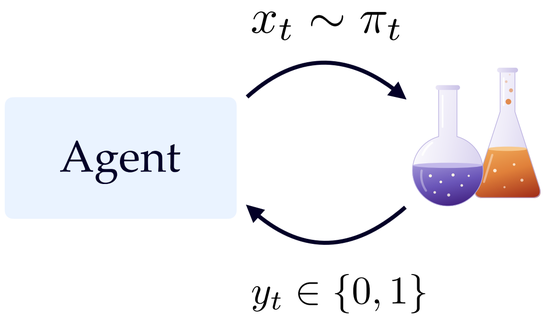
Black-box noisy binary feedback



Generable Set Expansion: Beyond Distribution Matching



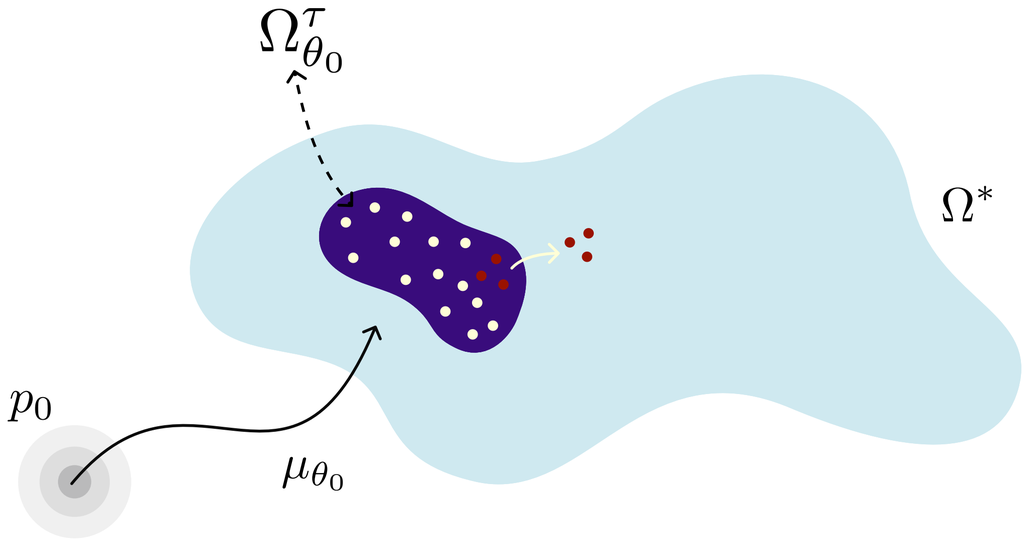
Black-box noisy binary feedback



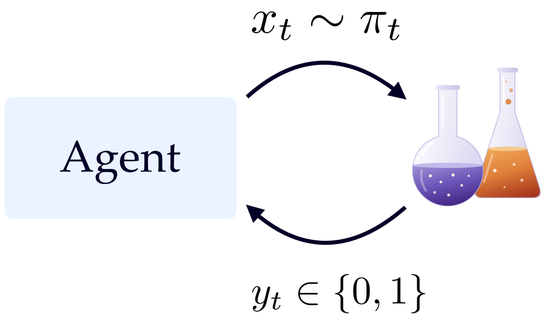
$\theta^{\text{pre}} = \theta_0$ $\Omega^{\tau}_{\theta_0}$

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

Generable Set Expansion: Beyond Distribution Matching



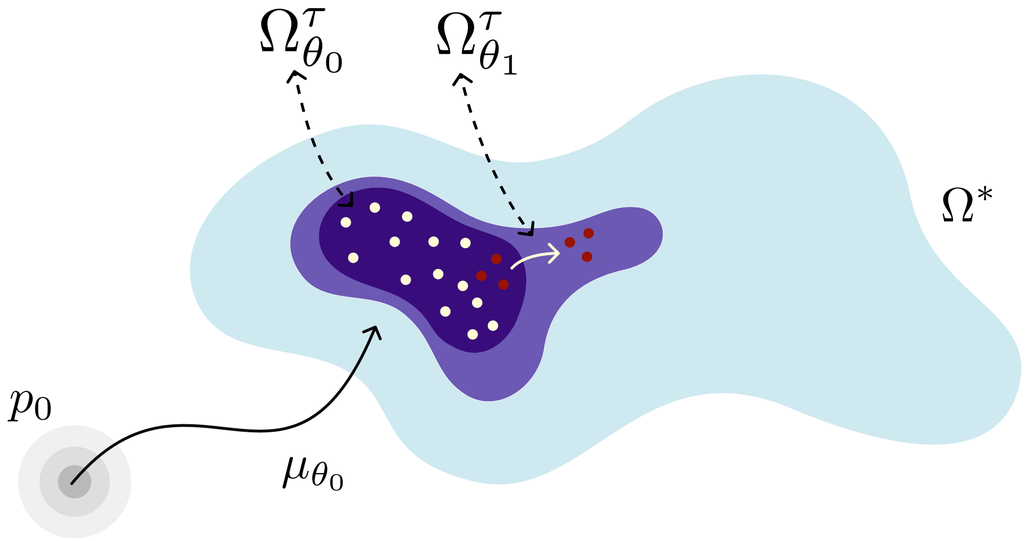
Black-box noisy binary feedback



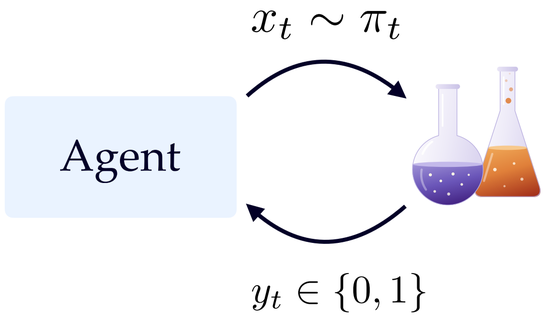
$\theta^{\text{pre}} = \theta_0$ $\Omega^{\tau}_{\theta_0}$

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

Generable Set Expansion: Beyond Distribution Matching



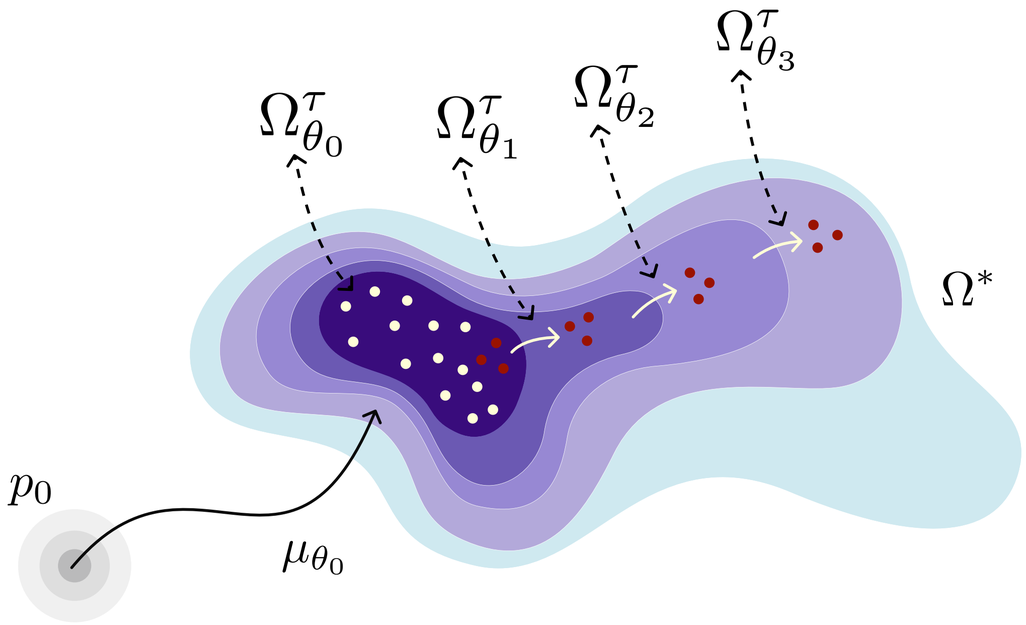
Black-box noisy binary feedback



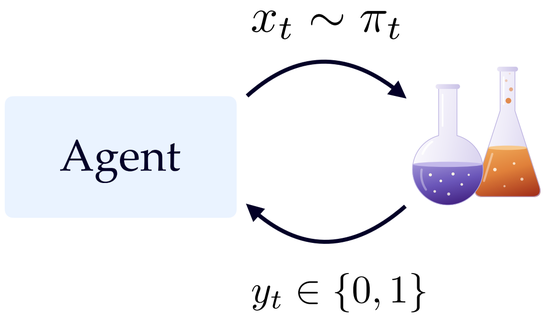
$$\theta^{\text{pre}} = \theta_0 \rightarrow \theta_1 \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1}$$

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

Generable Set Expansion: Beyond Distribution Matching



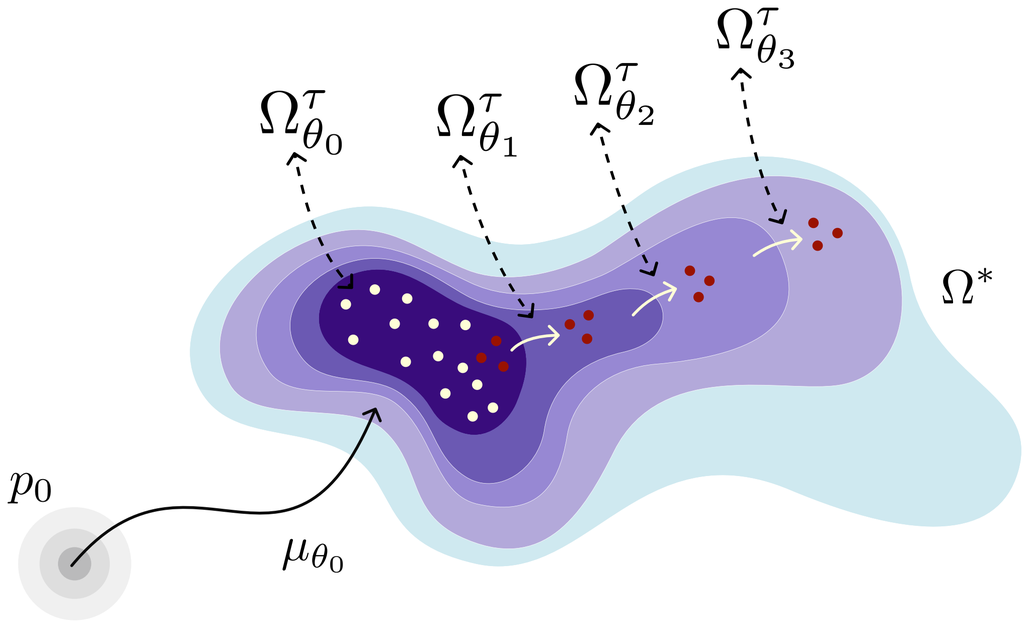
Black-box noisy binary feedback



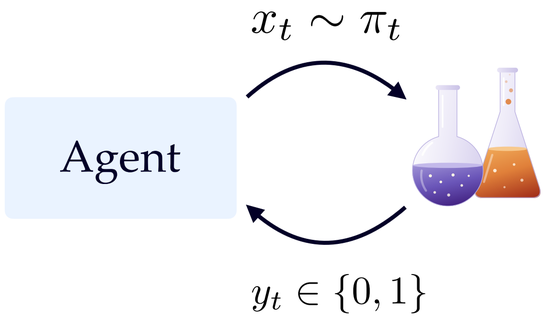
$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

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

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



Black-box noisy binary feedback



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 control for **active data collection.**

Statistical Guarantees for **Lab-in-the-Loop Active Generative Modeling**

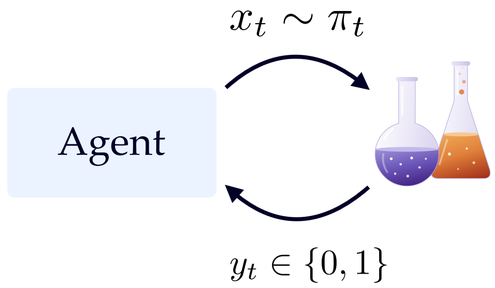
Distribution matching

$$p_1^\theta \approx p_{data}$$

Our work →

Design space coverage

$$\Omega_\theta^\tau \approx \Omega^*$$



Theorem (Informal). ActFlow guarantees, with probability at least $1 - \delta$:

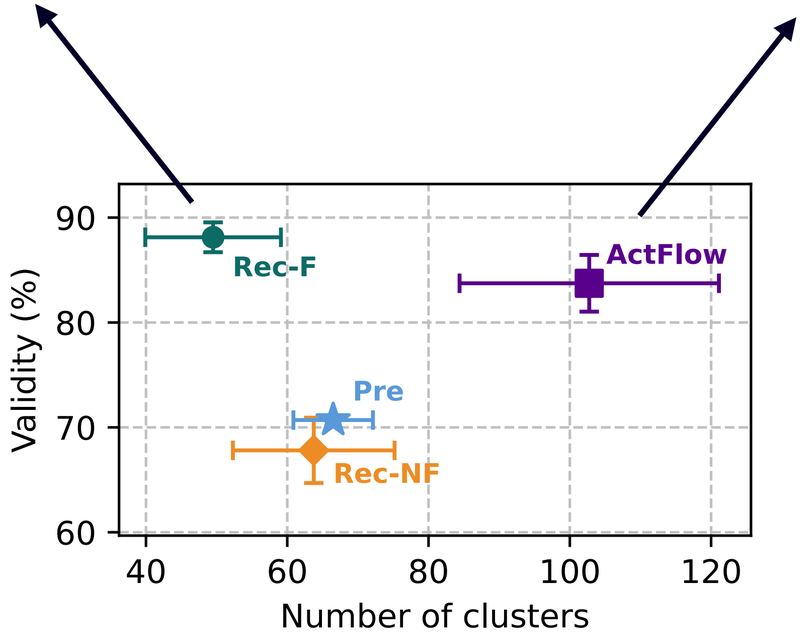
$$T^\star \gtrsim \left(\frac{\alpha \gamma_{HT^*}^{\Omega_\star}}{\epsilon} \right)^2 \implies R_\epsilon^H(S_0) \subseteq \Omega_{T^\star}^\tau$$

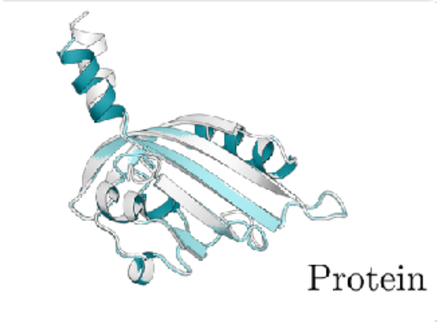
After sufficient verifier queries, the **generable set covers** the **H-step reachable set**.

ActFlow on Proteins

Model collapse!

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





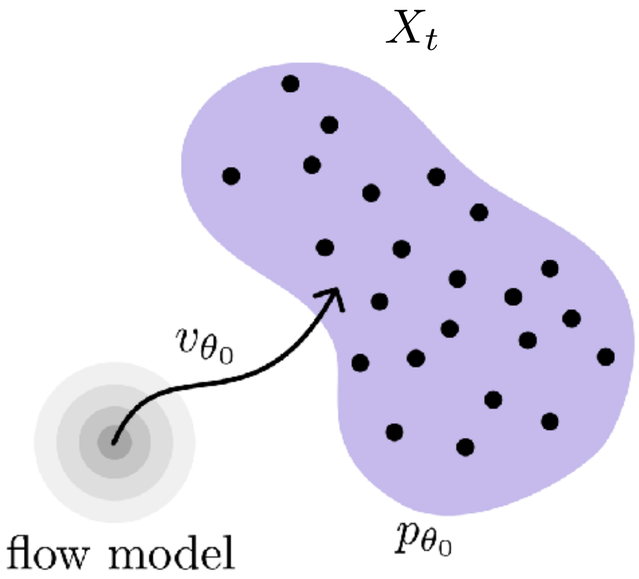
Sample-Efficient Goal-Directed Expansion for Discovery?

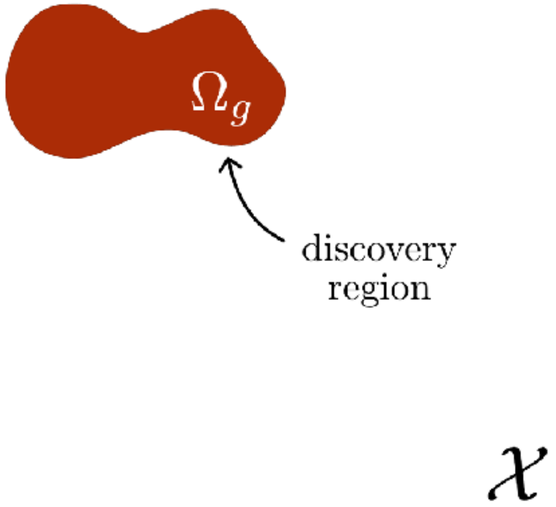
Natural-language discovery goal

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

Discovery region

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





LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

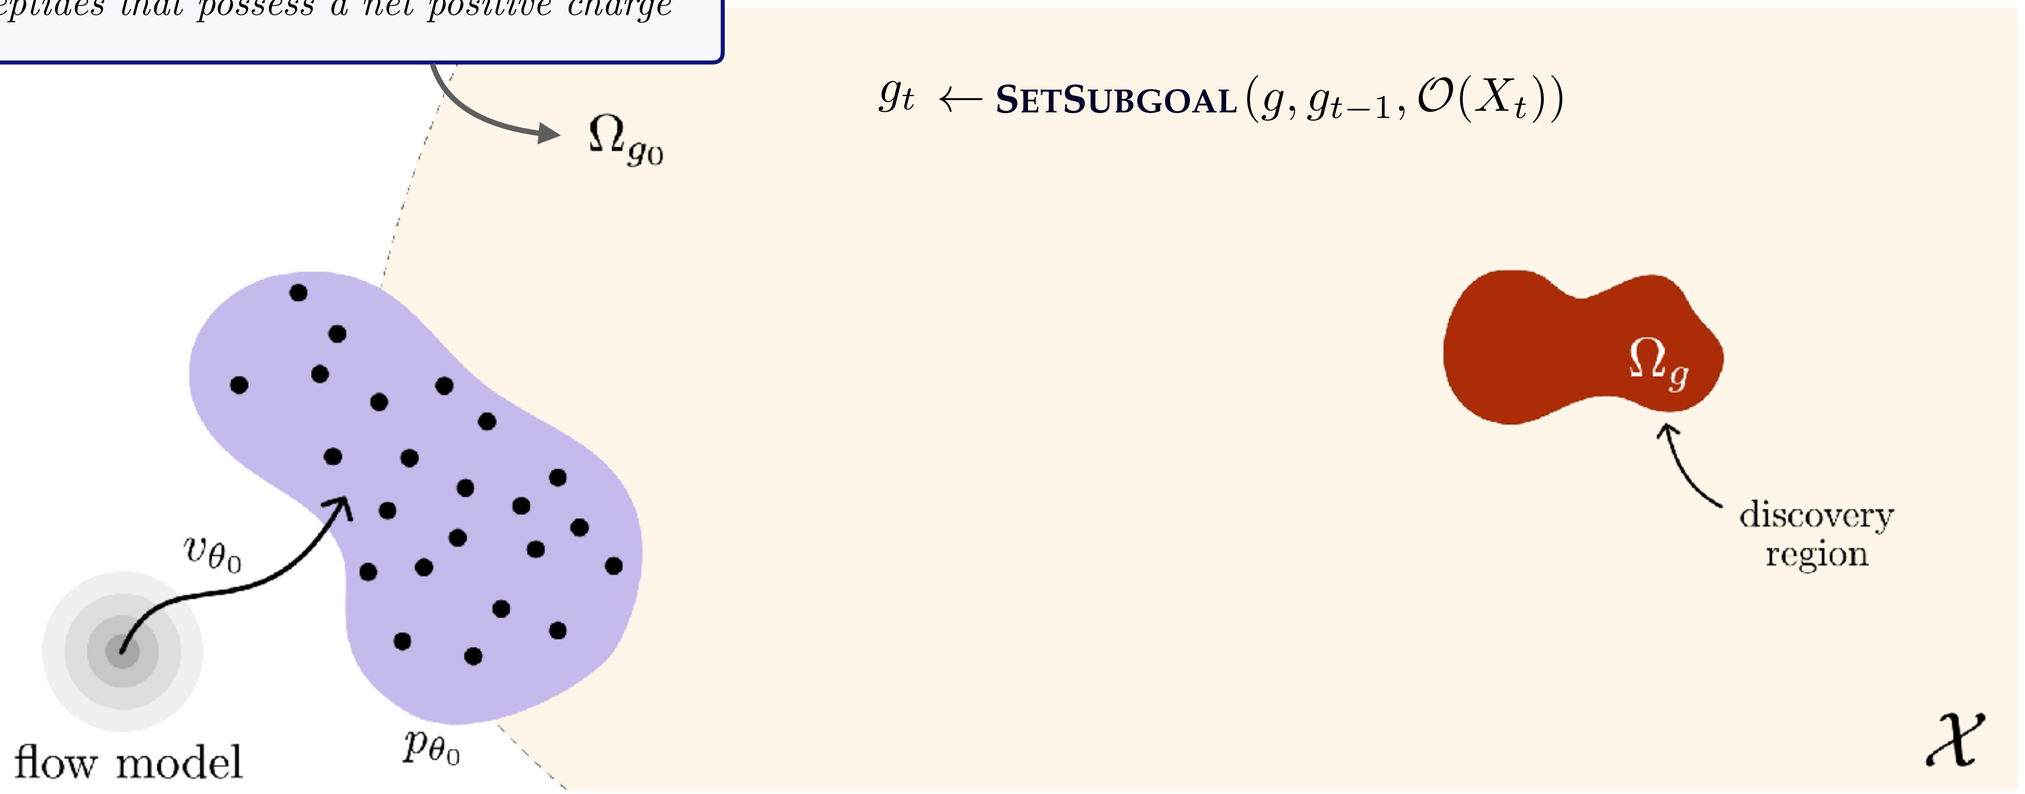
$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

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

$g_0 =$ *'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'*

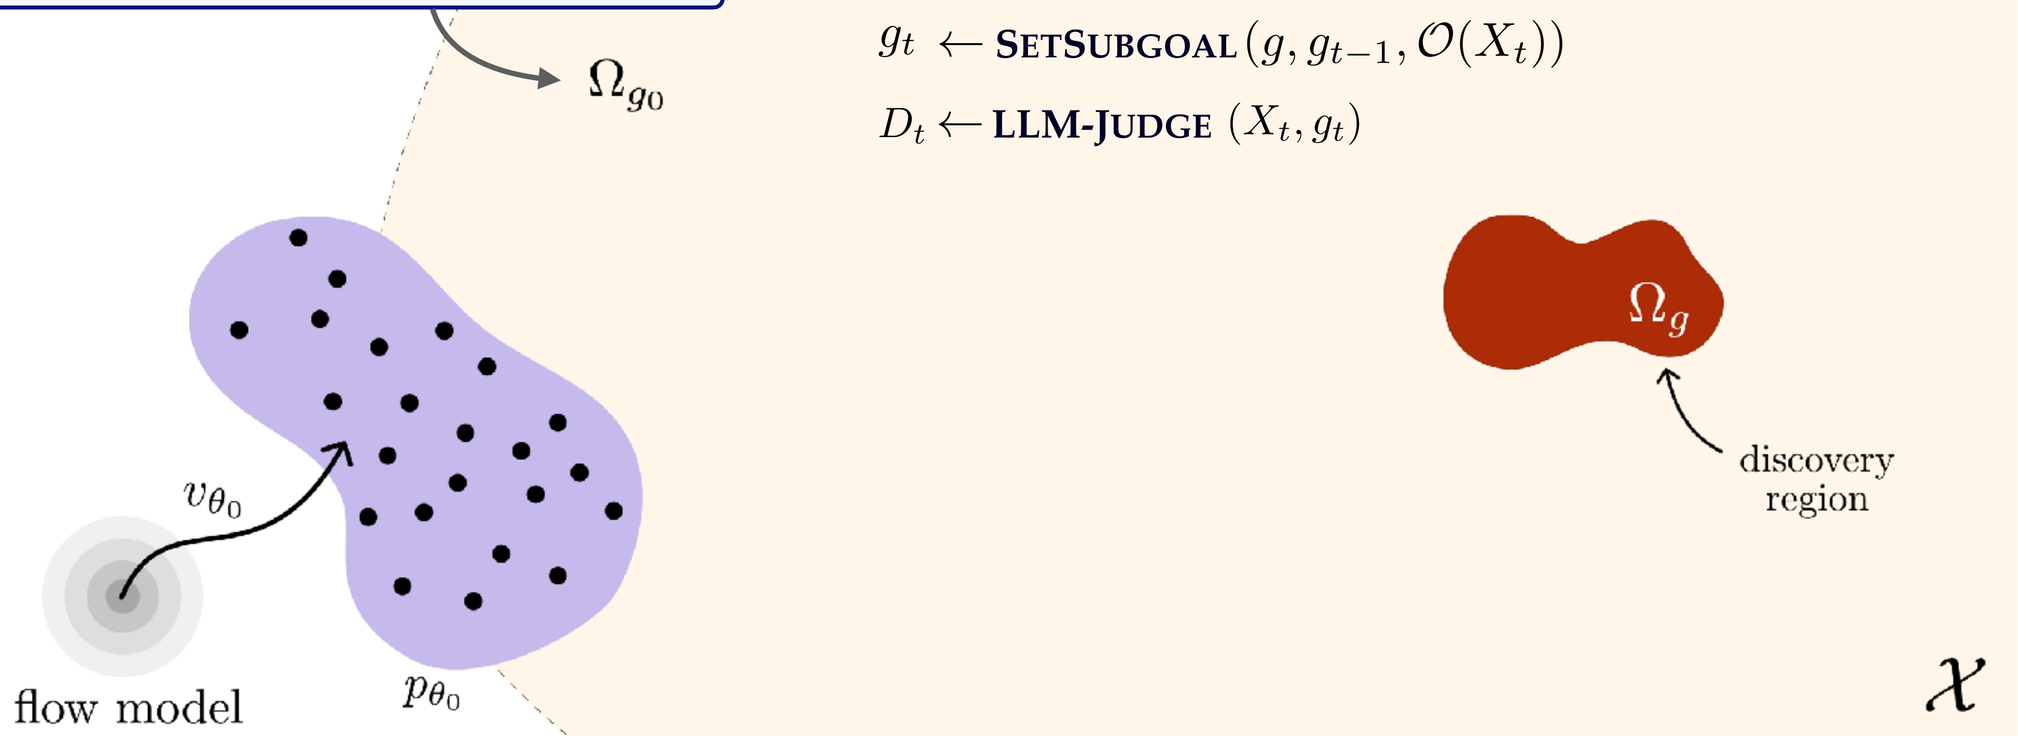
Discovery region

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

$g_0 =$ *'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)

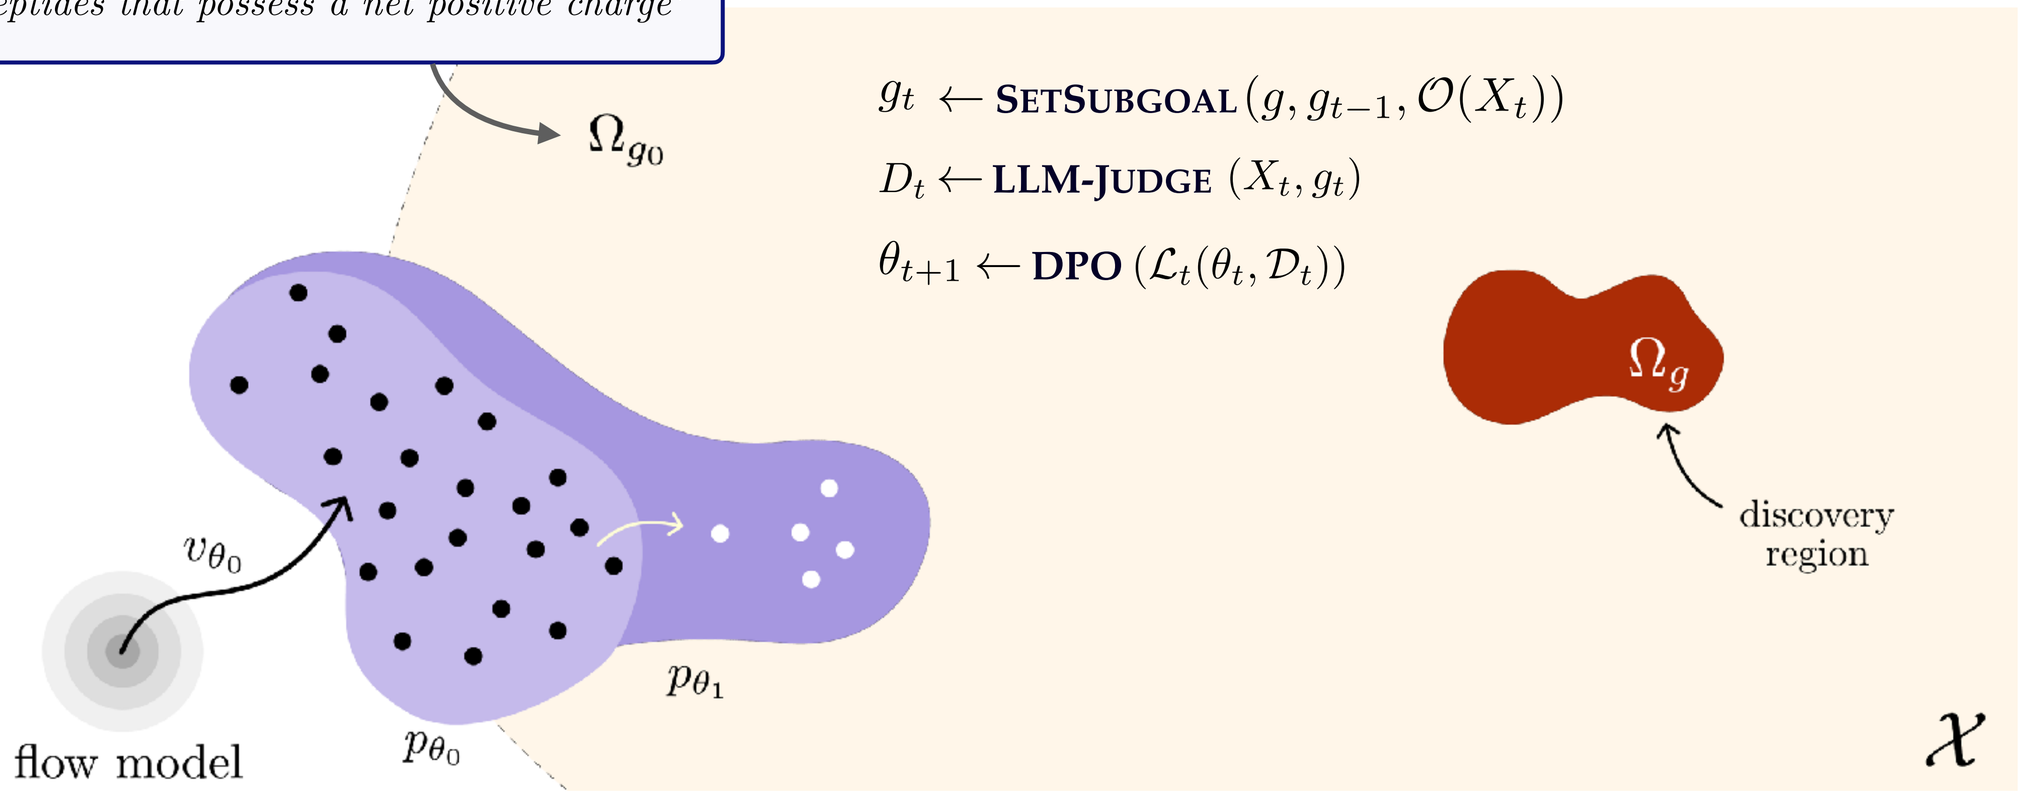
Natural-language discovery goal

$g =$ *'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'*



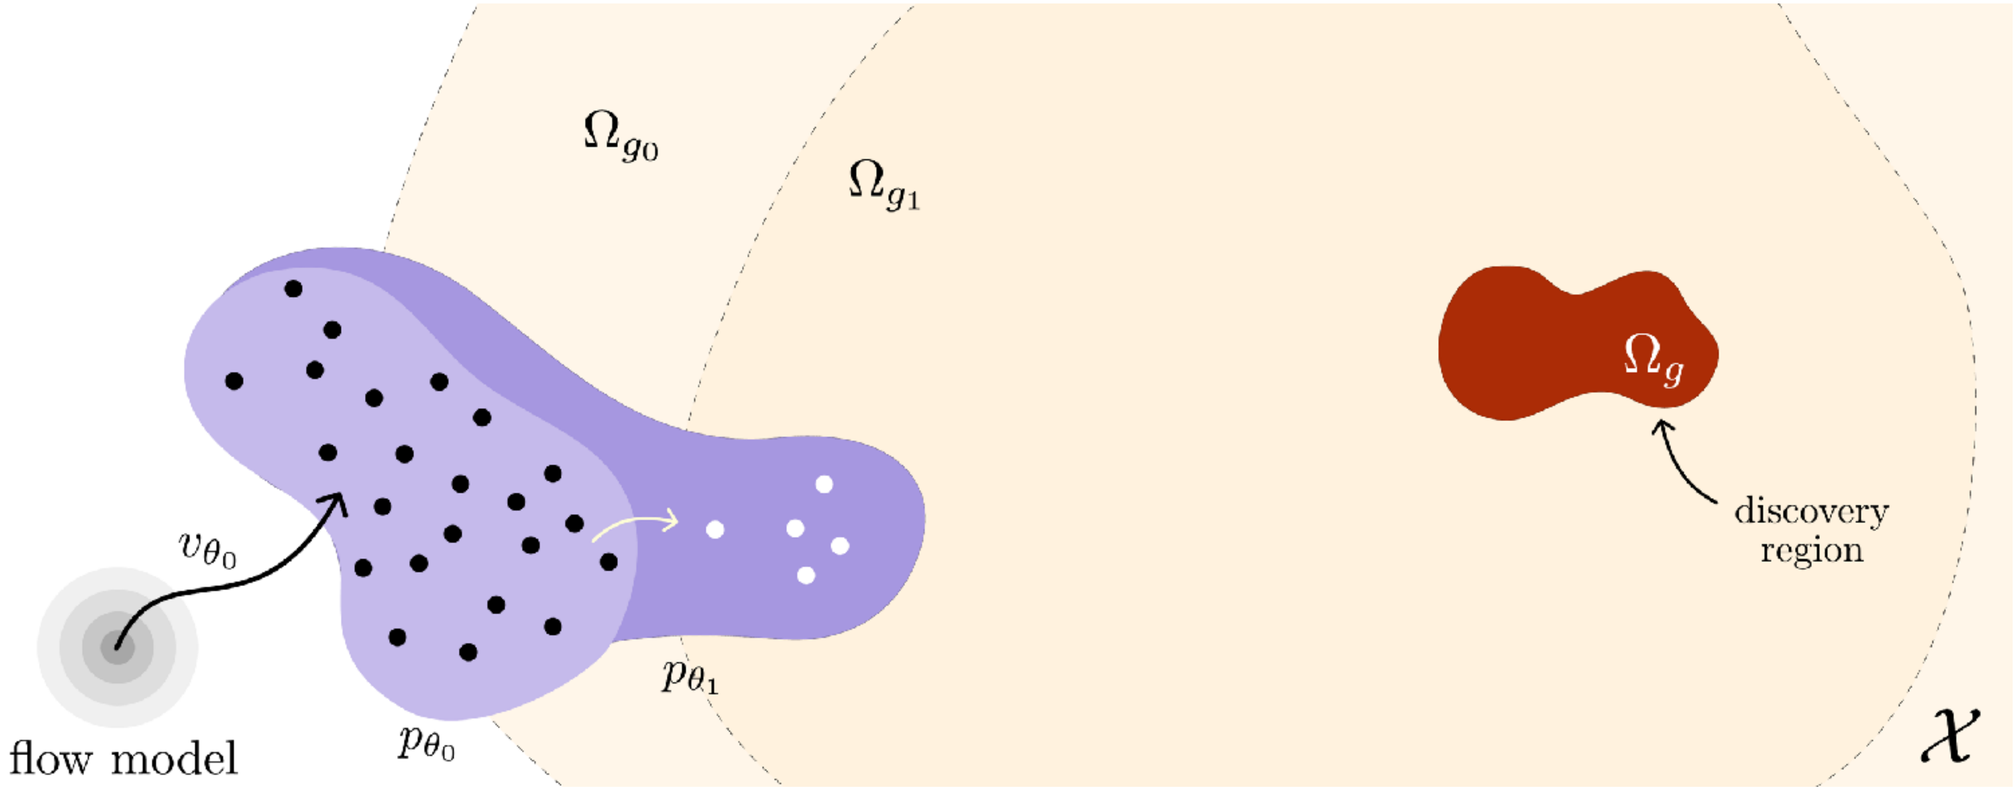
LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

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

Discovery region

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



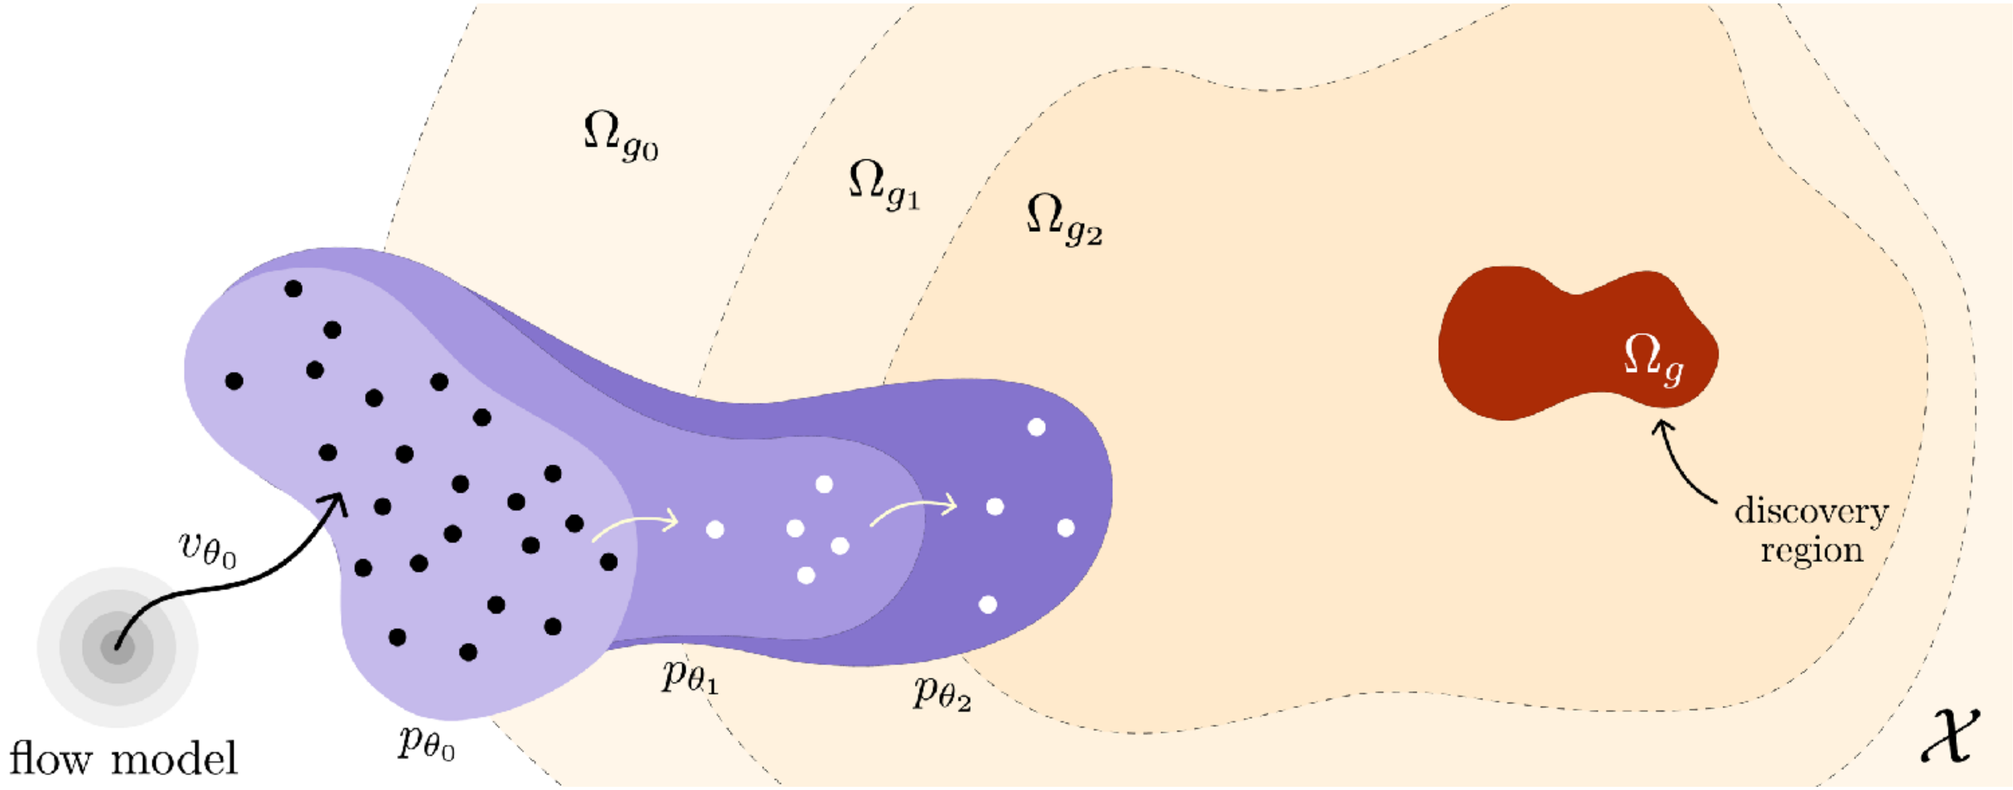
LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

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

Discovery region

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



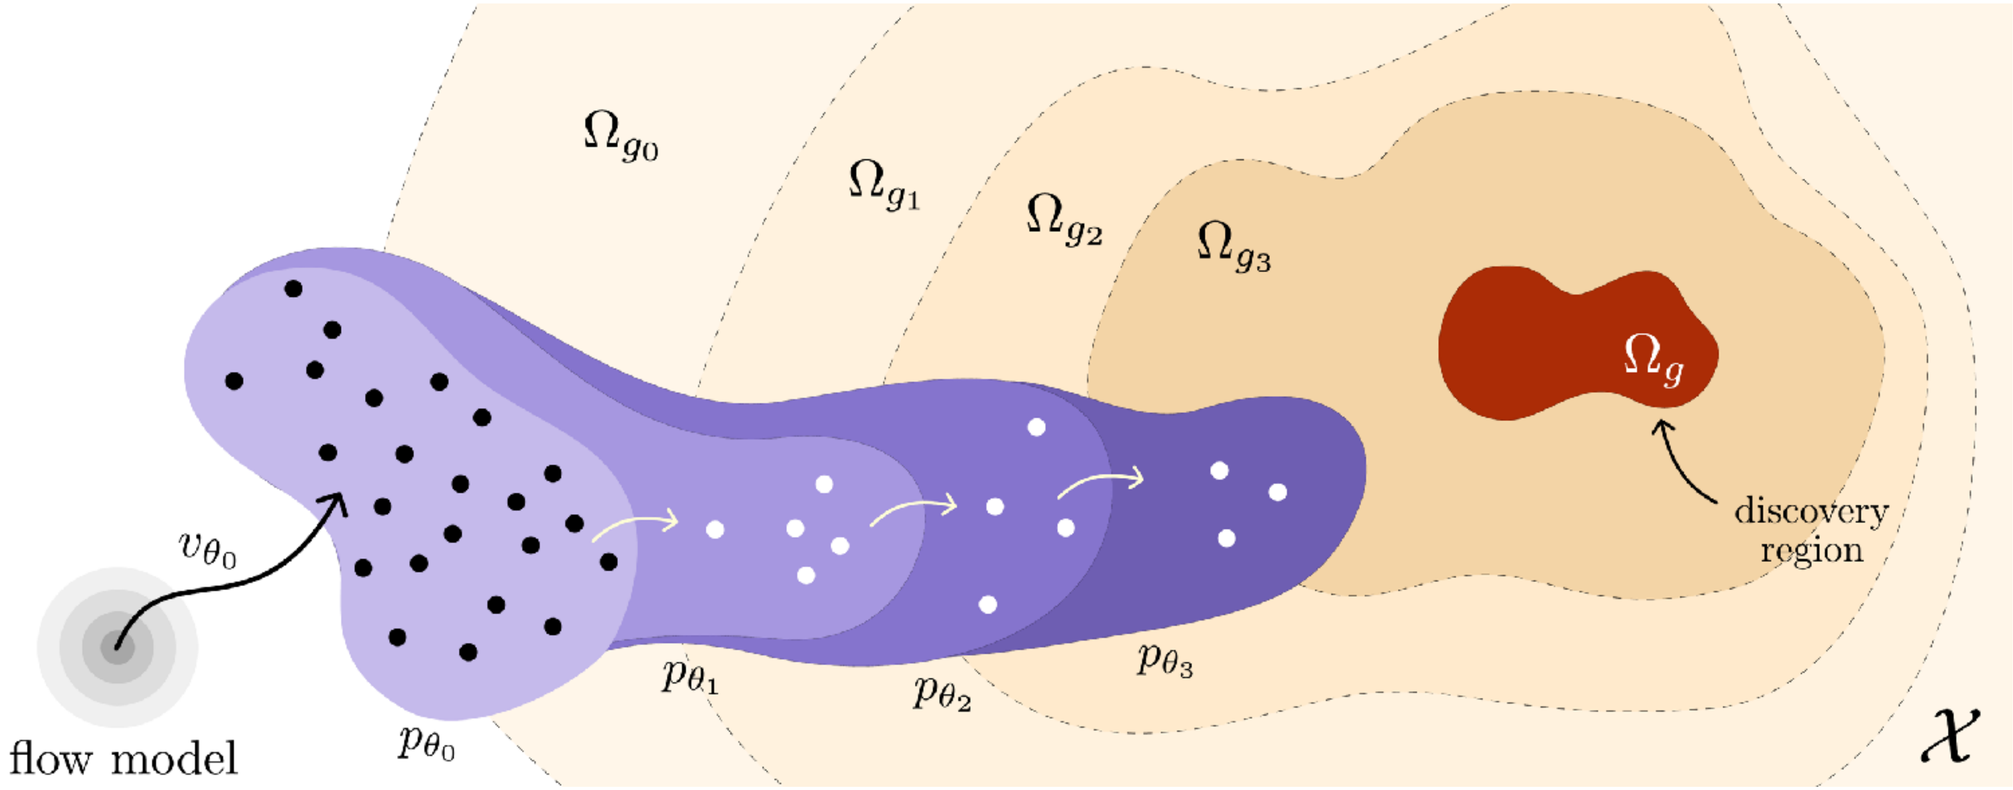
LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

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

Discovery region

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



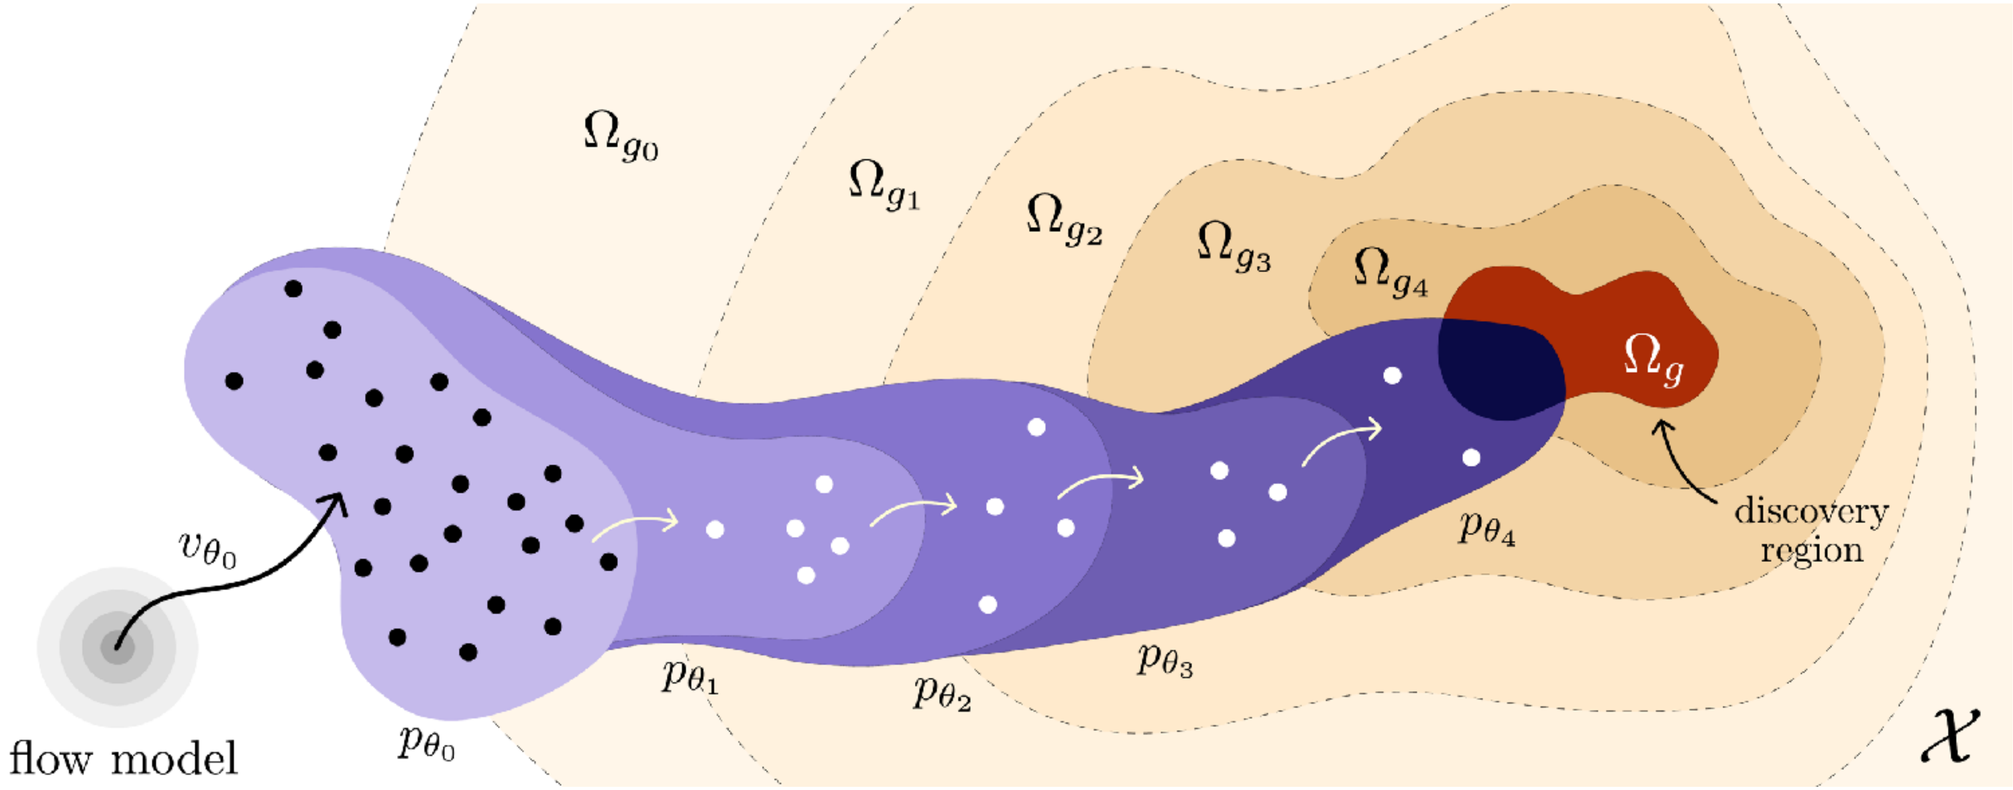
LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

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

Discovery region

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



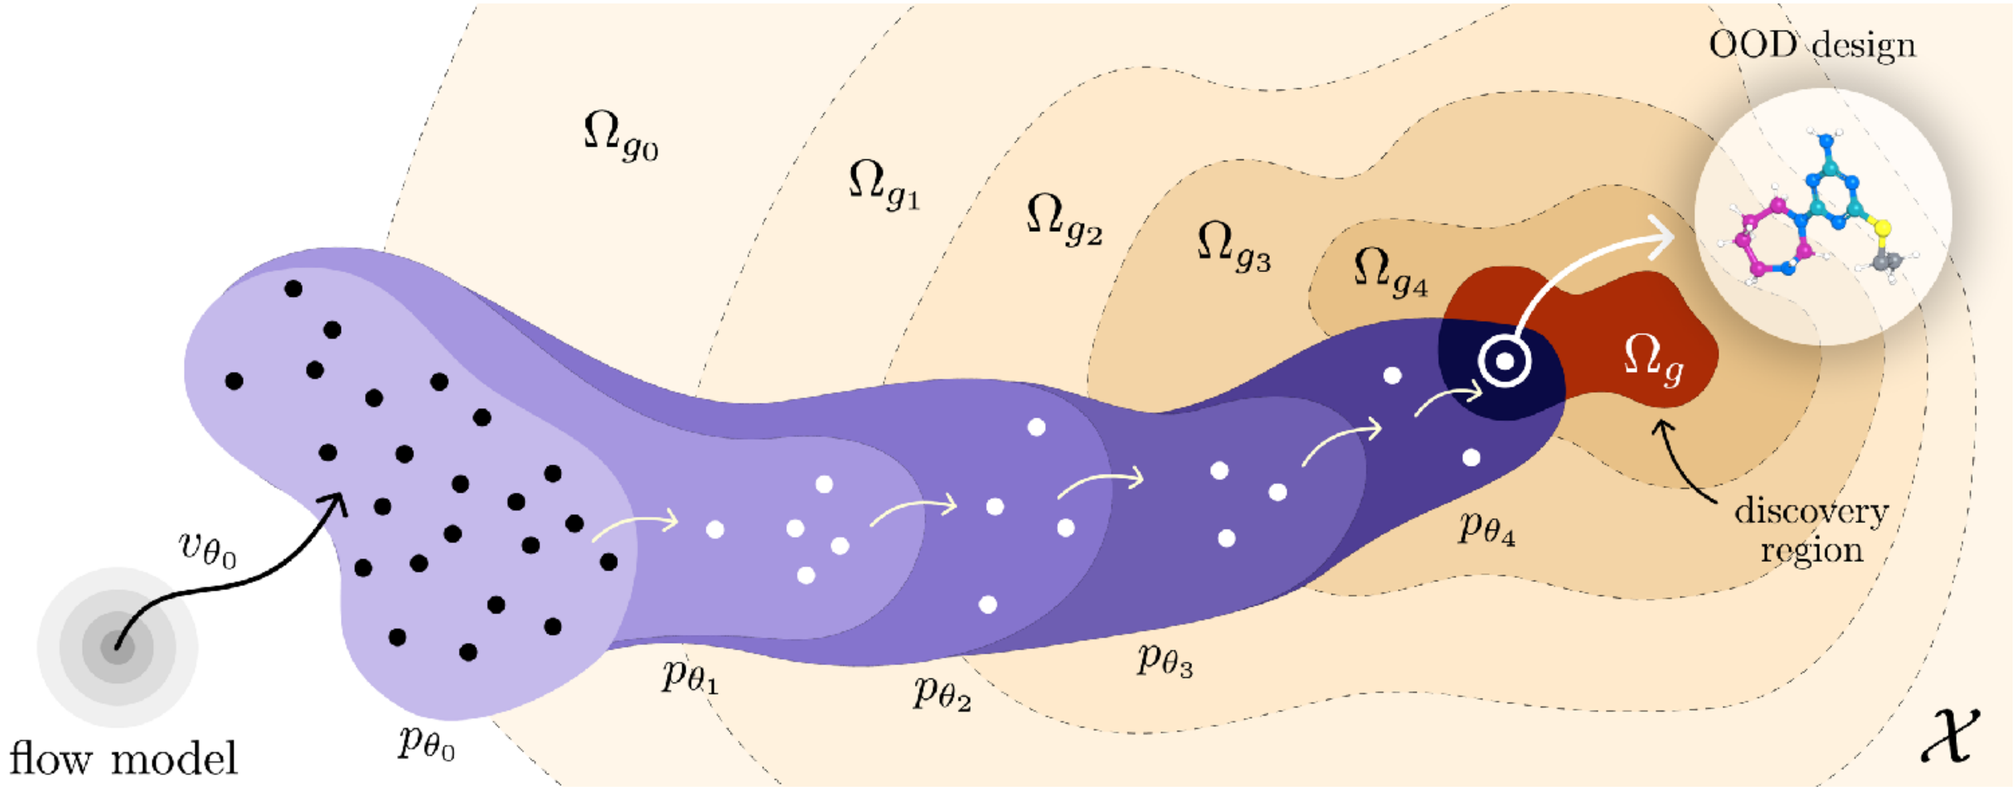
LLM-to-Flow Distillation (LFD)

Natural-language discovery goal

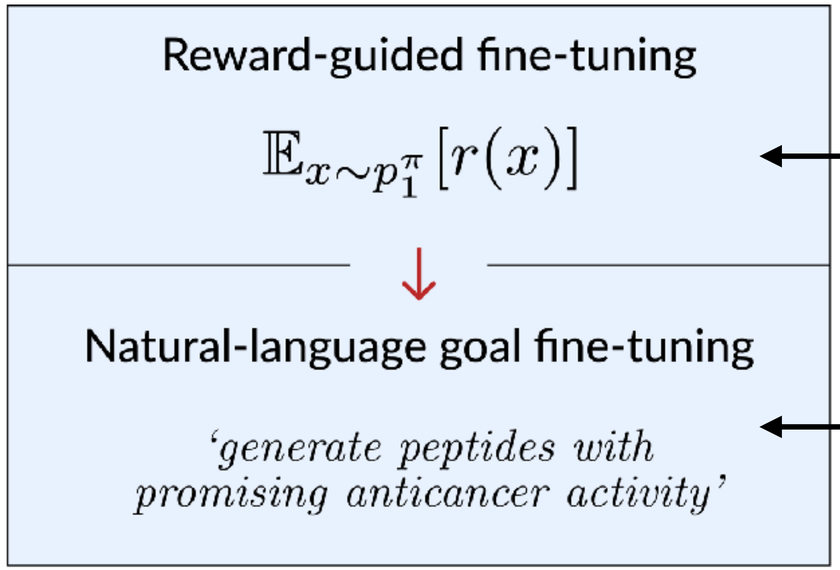
$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

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

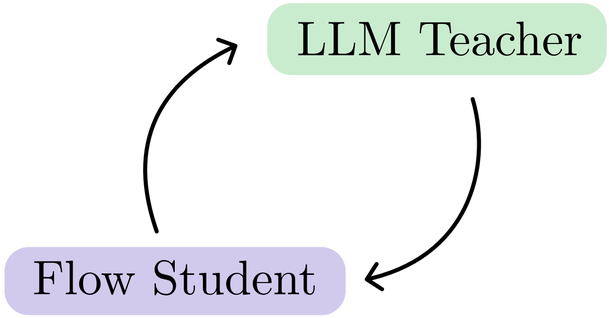


LFD: From Rewards to Natural-Language Discovery Goals



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**

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.

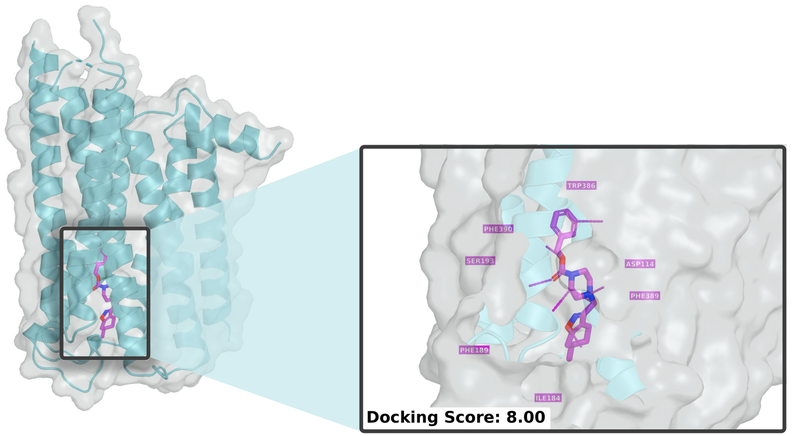
53.5% (LFD) vs 35.9% (best reward-guided baseline) on a held-out anticancer predictor.

GOAL 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**.

↓

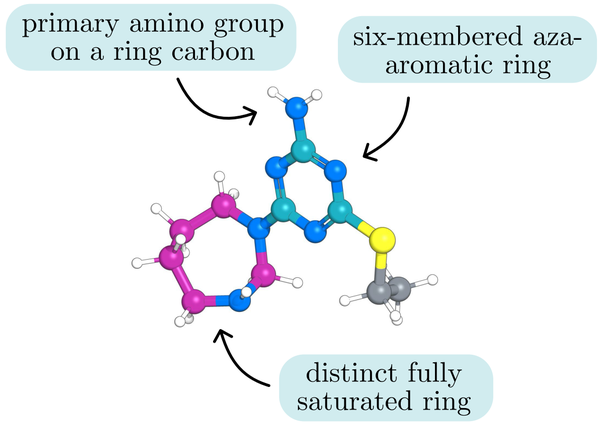
Desired structure is **out of** prior model's **generable set**

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.



Mathematical Foundations of Exploration

Exploration in Reinforcement Learning

*"The Importance of Non-Markovianity in Maximum State **Entropy** Exploration"*
Outstanding Paper Award at **ICML 2022**

*"Challenging Common Assumptions in **Convex** Reinforcement Learning"*
NeurIPS 2022

*"**Geometric** Active Exploration in Markov Decision Processes: the Benefit of Abstraction"*
ICML 2024

*"Exploiting **Causal** Graph Priors with Posterior Sampling for Reinforcement Learning"*
ICLR 2024

*"Global Reinforcement Learning: Beyond Linear and Convex Rewards via **Submodular** Semi-gradient Methods"*
ICLR 2024

Mathematical Foundations of Exploration

↓

Generative Discovery Algorithms

→

Generative Discovery via Diffusion Models Adaptation

"Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning"
Spotlight at **NeurIPS 2025, Oral at Generative AI and Biology Workshop**

"Provable Maximum Entropy Manifold Exploration via Diffusion Models"
ICML 2025

"Verifier-Constrained Flow Expansion for Discovery Beyond the Data"
ICLR 2026

"Constrained Flow Optimization via Sequential Fine Tuning for Molecular Design"
ICLR 2026, Oral at Frontiers in Probabilistic Inference Workshop

"A Unified Density Operator View of Flow Control and Merging"
ICLR 2026, Oral at Real-World Constrained Diffusion-Based Models Workshop

"Active Flow Expansion for Out-of-Distribution Discovery: from Theory to Molecules"
NeurIPS 2026

"LLM-to-Flow Distillation: Teaching Natural Language Goals to Scientific Generators"
Preprint.

Mathematical Foundations of Exploration

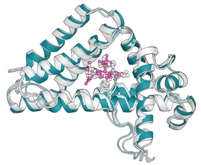
↓

Generative Discovery Algorithms

↓

Experimental Sciences (Wet-lab Discovery)

Enzymes

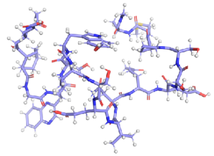




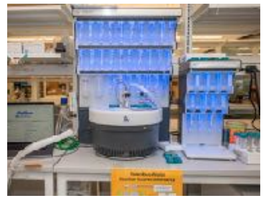


Caltech

Therapeutic Peptides

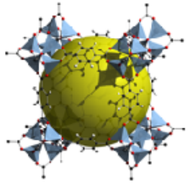






Penn
University of Pennsylvania

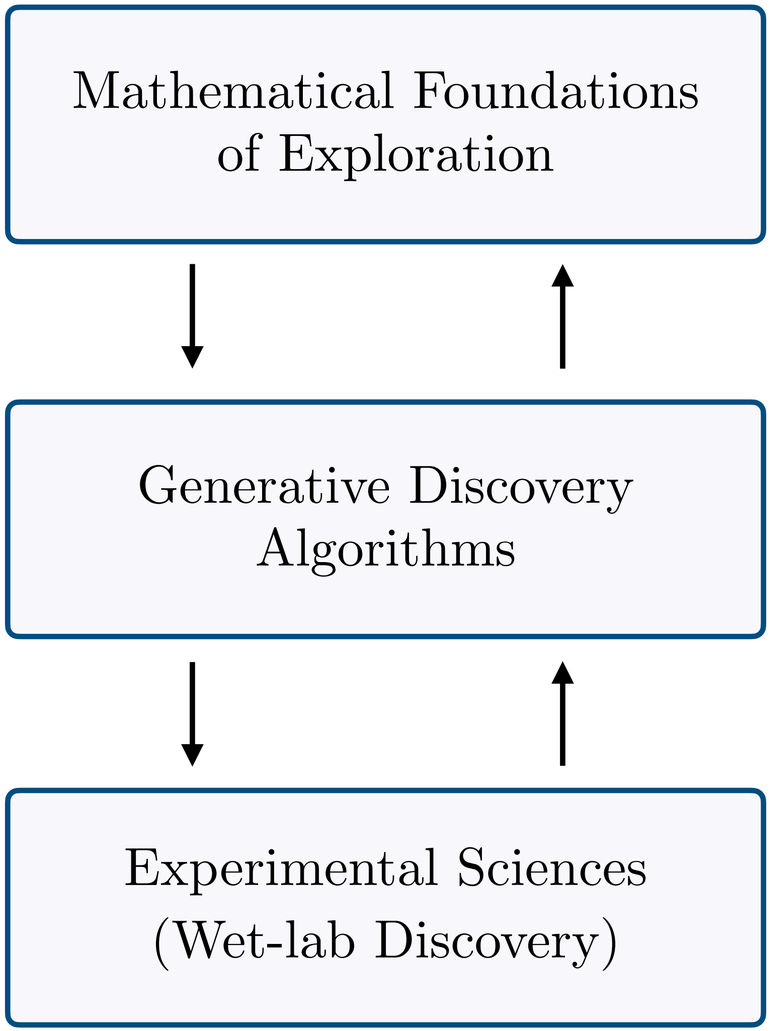
Functional Materials



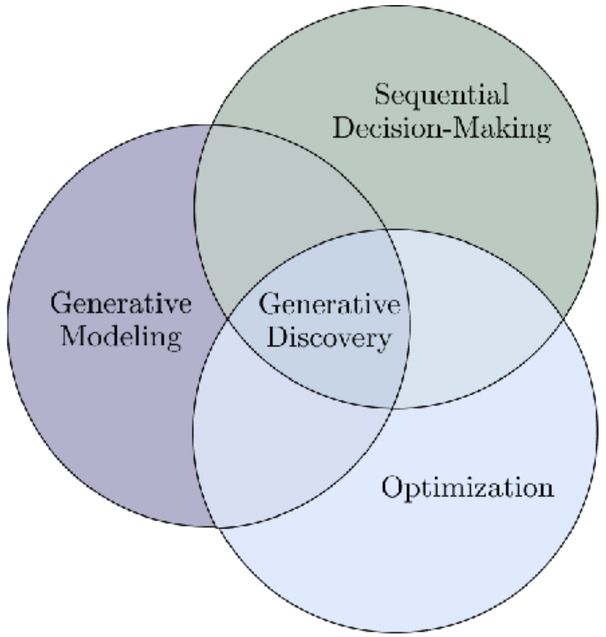




ETH zürich



Toward a **science of generative discovery** for **experimental sciences**, with **exploration as its unifying principle**.



Thanks!

Collaborators

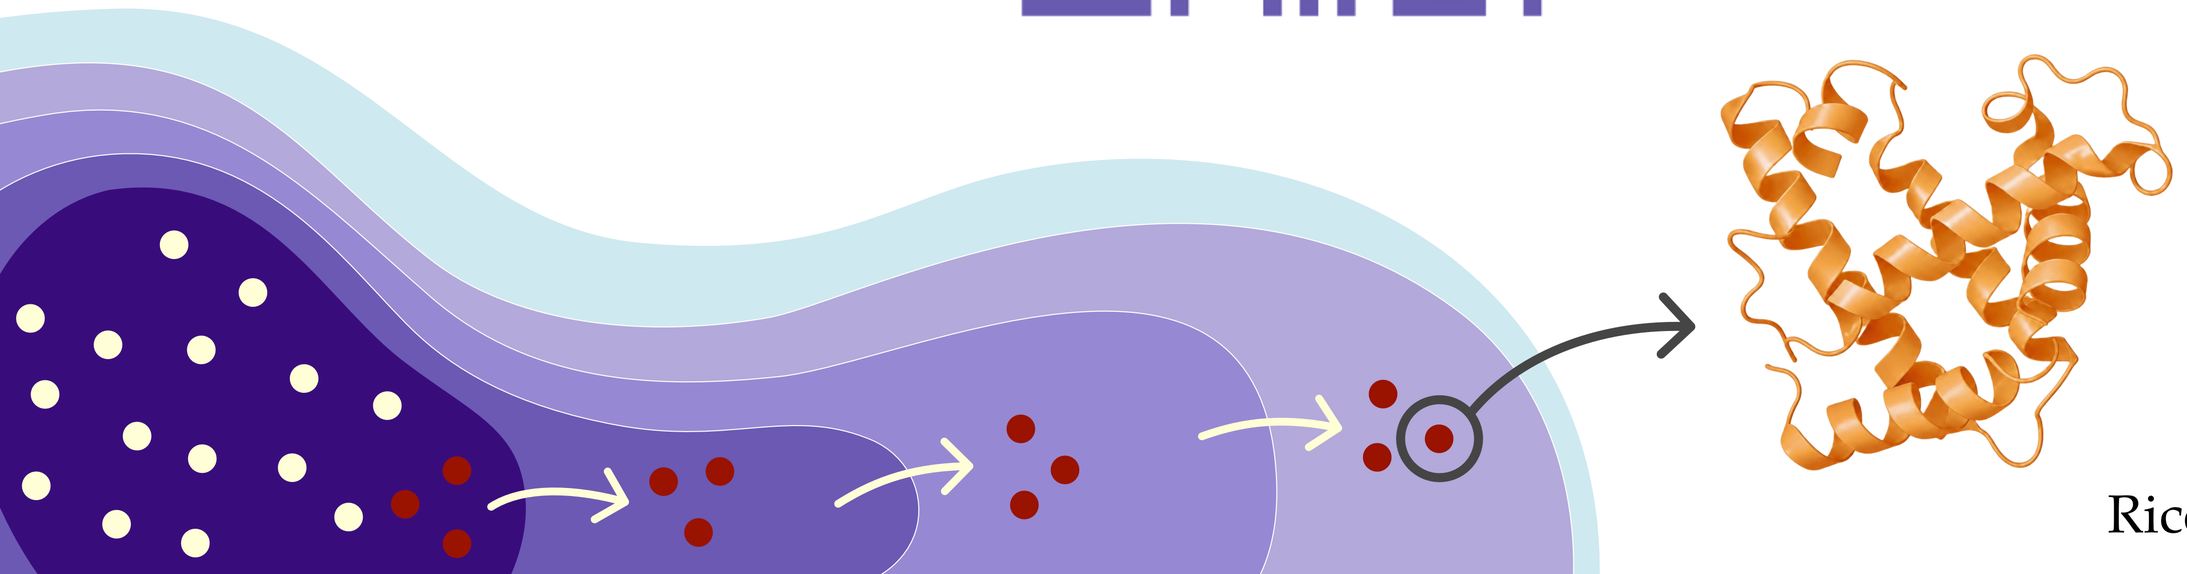
Kimon Protopapas, Ya-Ping Hsieh, Sven Gutjahr, Luca Schaufelberger, Zifan Wang, Xiaoyu Mo, Malte Franke, Cristian Perez Jensen, Andrey Kharitenko, Zebang Shen, Marin Vlastelica, Guy Schacht, Mojmir Mutny, Ziyad Sheebaelhamd, Bruce Lee, Sophia Tang, Cheng-Hao Liu, Federico Di Gennaro, Angelo Gnazzo, Michael 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.

Thanks!

My Research

riccardodesanti.com/research

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



Riccardo De Santi