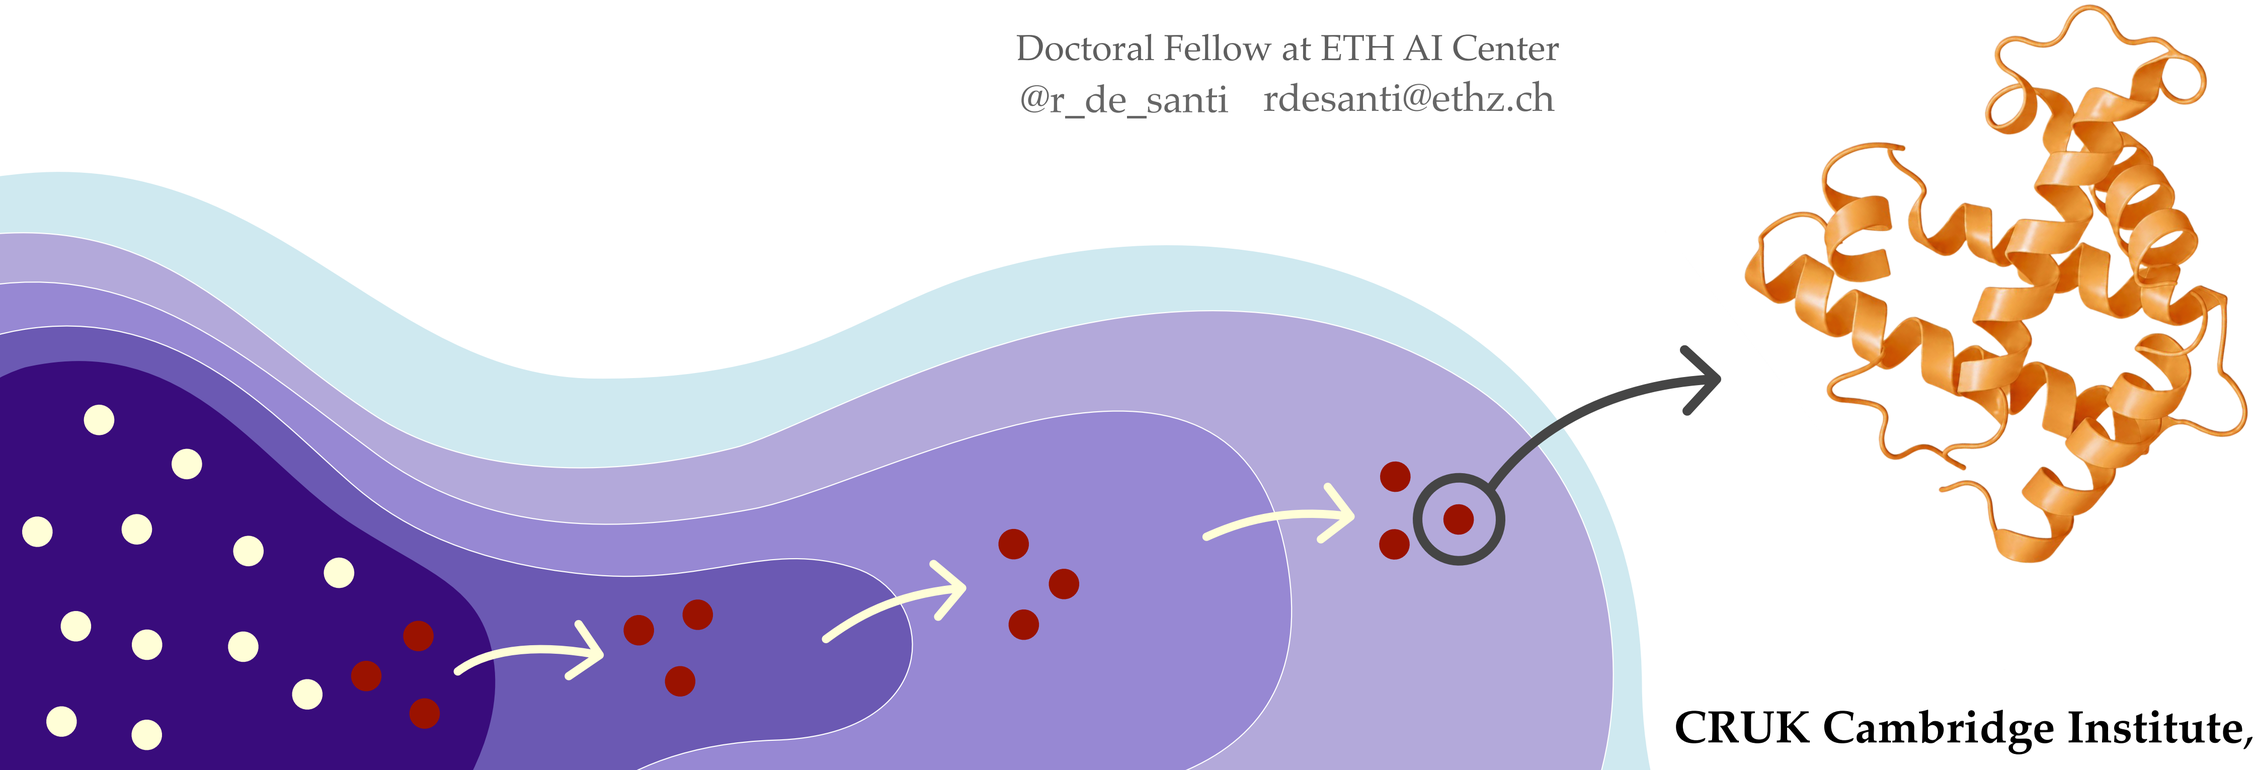
Generative Discovery Beyond the Data

Riccardo De Santi

Doctoral Fellow at ETH AI Center

@r_de_santi rdesanti@ethz.ch



CRUK Cambridge Institute, Nov. 5th, 2026

Team



Riccardo De Santi



Ya-Ping Hsieh



Kimon Protopapas



Bruce Lee



Cristian P. Jensen



Zebang Shen



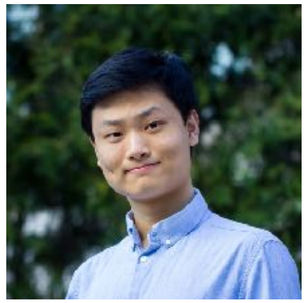
Marin Vlastelica



Zifan Wang



Sophia Tong



Chenghao Liu



Federico Di Gennaro



Angelo Gnazzo

Advisors


Pranam Chatterjee


Niao He


İsmail İlkan Ceylan


Michael M. Zavlanos


Fanny Yang

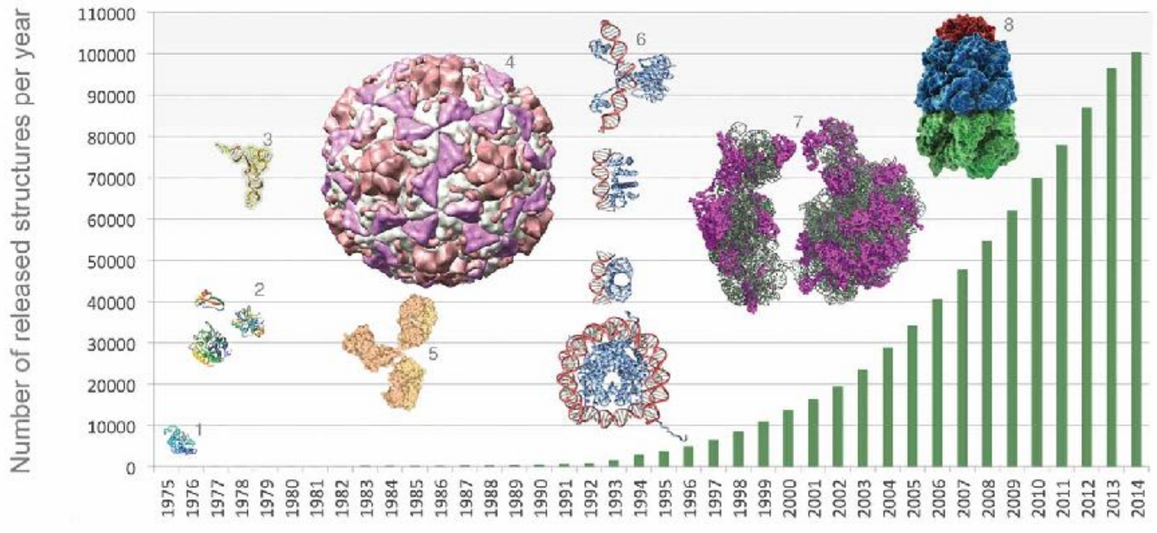

Karl H. Johansson


Frances H. Arnold

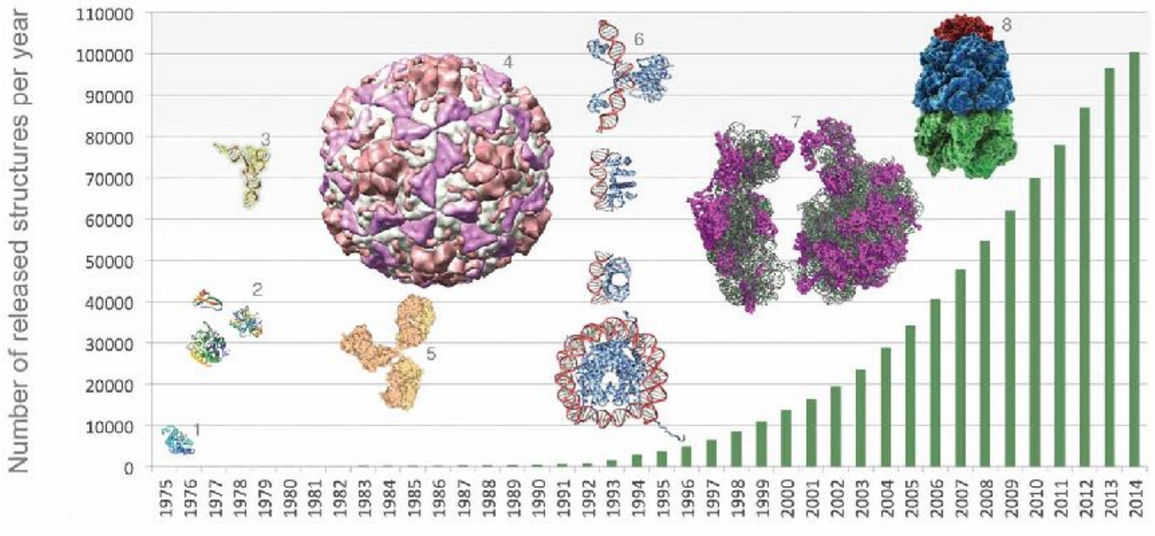

Yisong Yue


Michael M. Bronstein


Andreas Krause

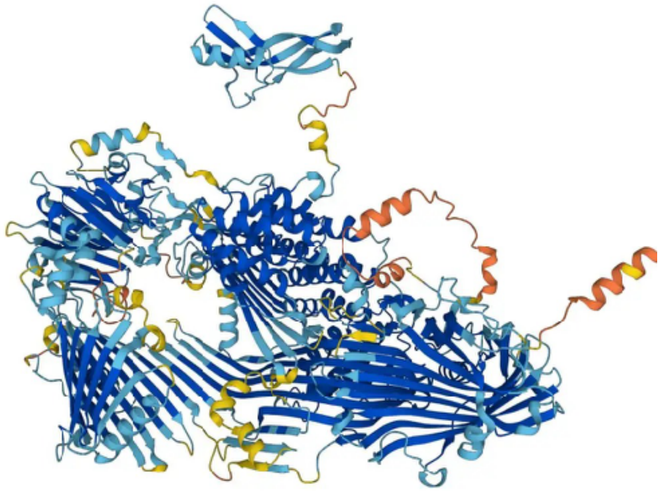


Protein Data Databank **(PDB)**

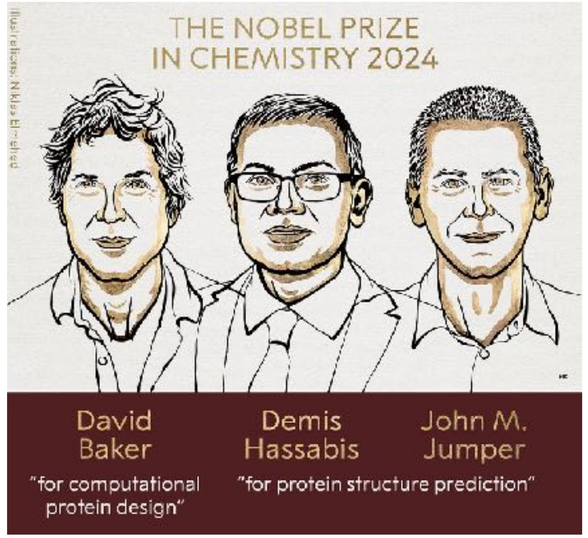


Protein Data Databank **(PDB)**

Protein structure prediction and design

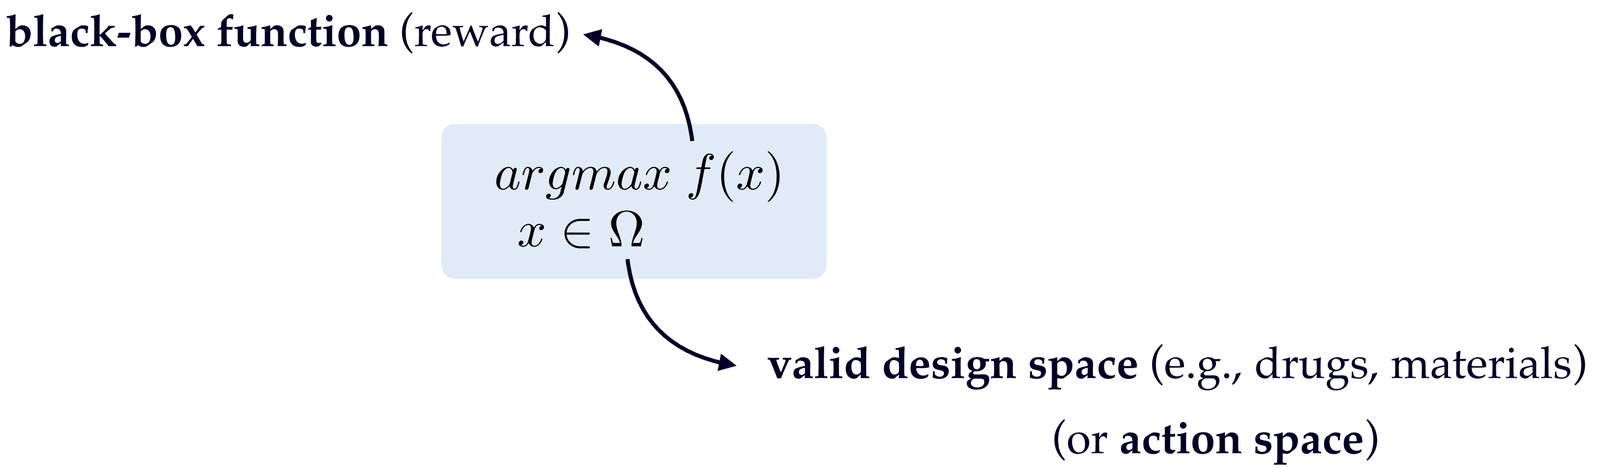


2024 Chemistry Nobel Prize



Generative ~~Modeling~~ Discovery?

What is Discovery?



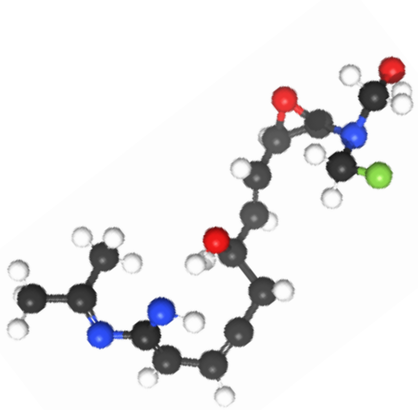
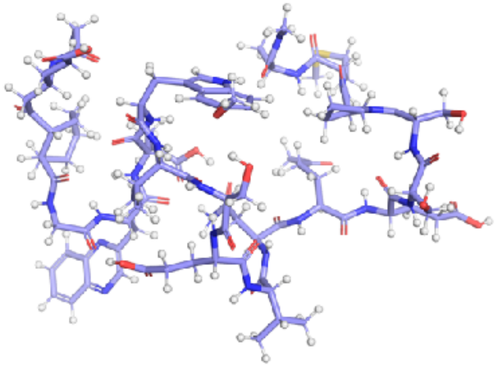
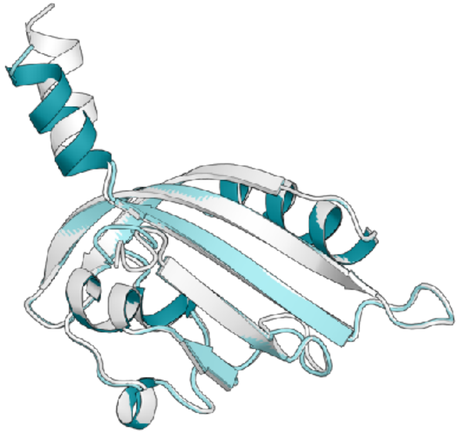
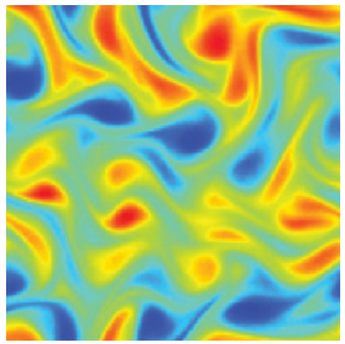
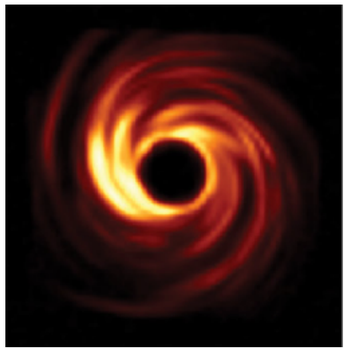
What is Discovery?

black-box function (reward)

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

valid design space (e.g., drugs, materials)

(or **action space**)

[Zheng et al., 2025]

What is Discovery?

black-box function (reward)

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

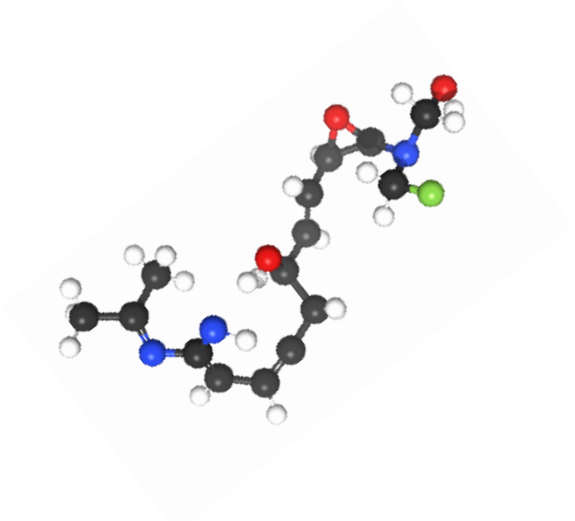
valid design space (e.g., drugs, materials)

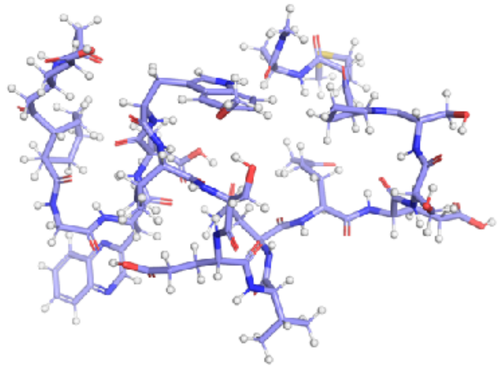
(or **action space**)

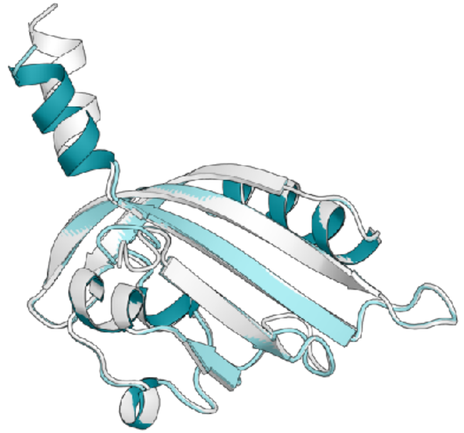
$\approx 10^{60}$ drug discovery

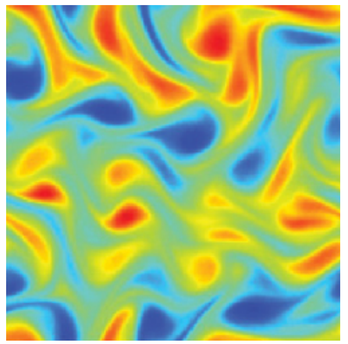
$\approx 10^{80}$ material design

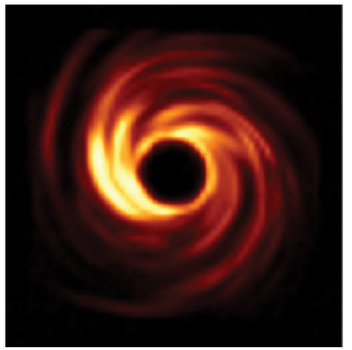
unknown or **hard to represent** and **search**











[Zheng et al., 2025]

What is Discovery?

black-box function (reward)

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

valid design space (e.g., drugs, materials)

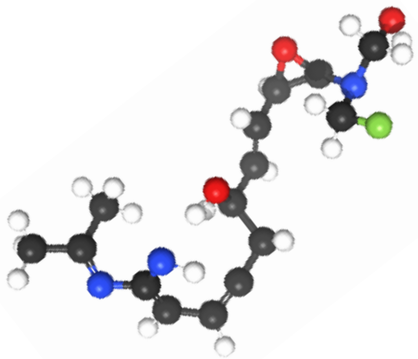
(or **action space**)

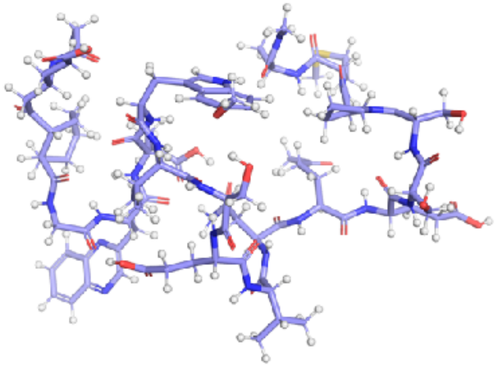
$\approx 10^{60}$ drug discovery

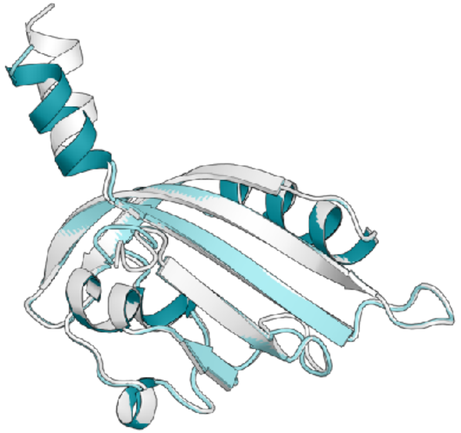
$\approx 10^{80}$ material design

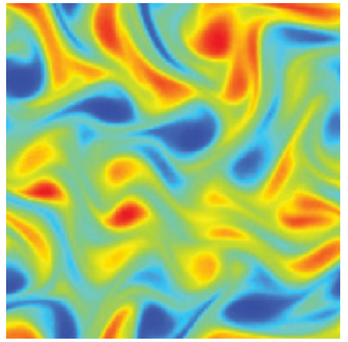
unknown or **hard to represent** and **search**

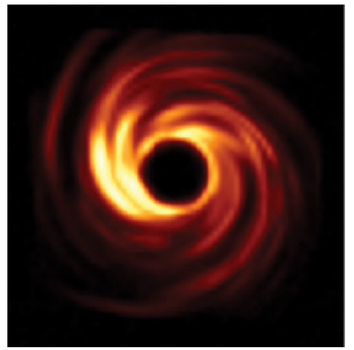
→ Generative Discovery







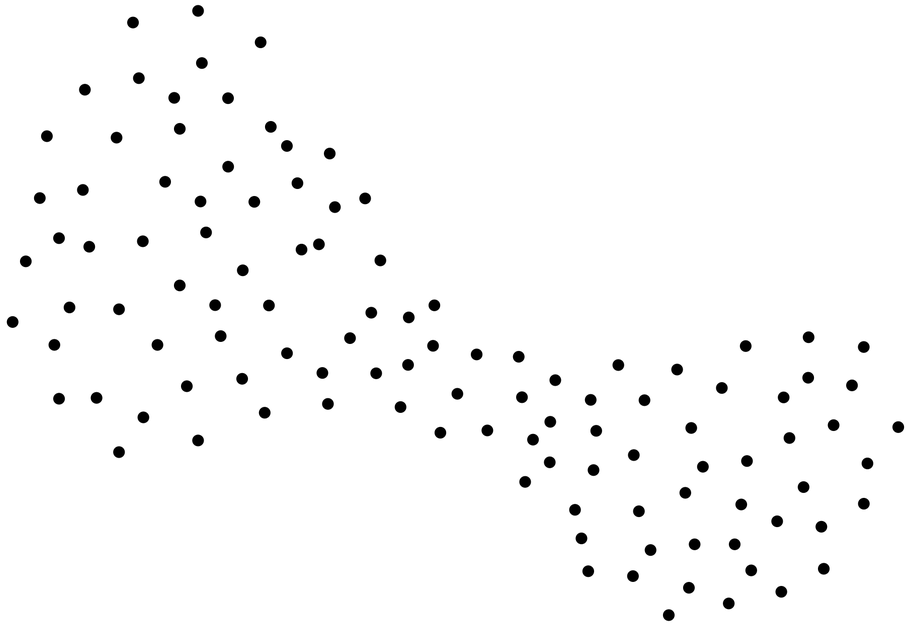




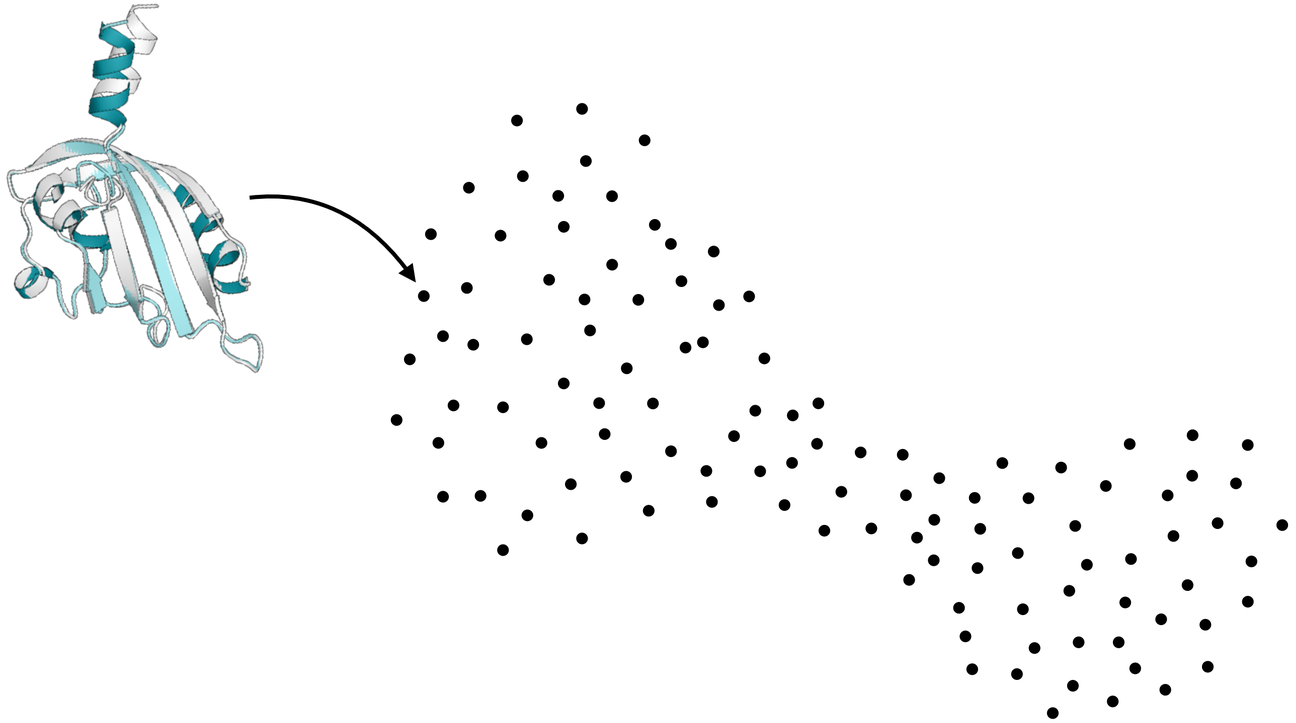
[Zheng et al., 2025]

Generative Discovery Paradigm

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



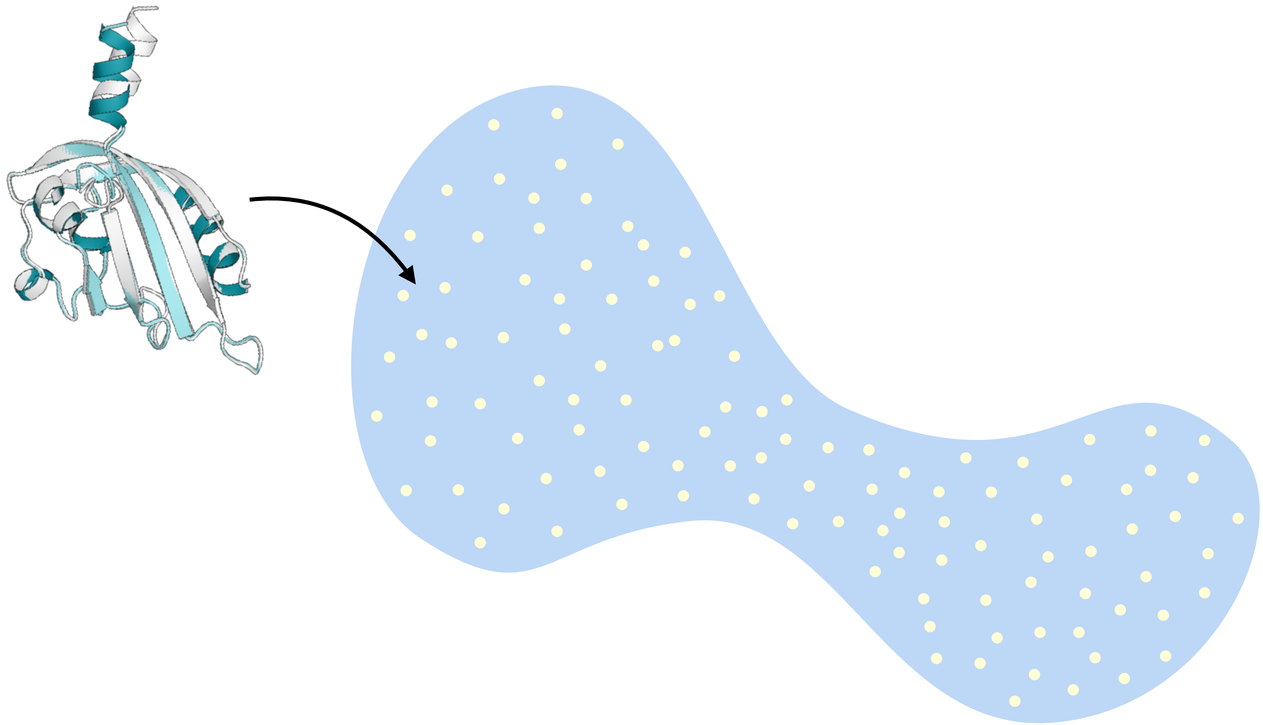
Generative Discovery Paradigm



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

Generative Discovery Paradigm

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

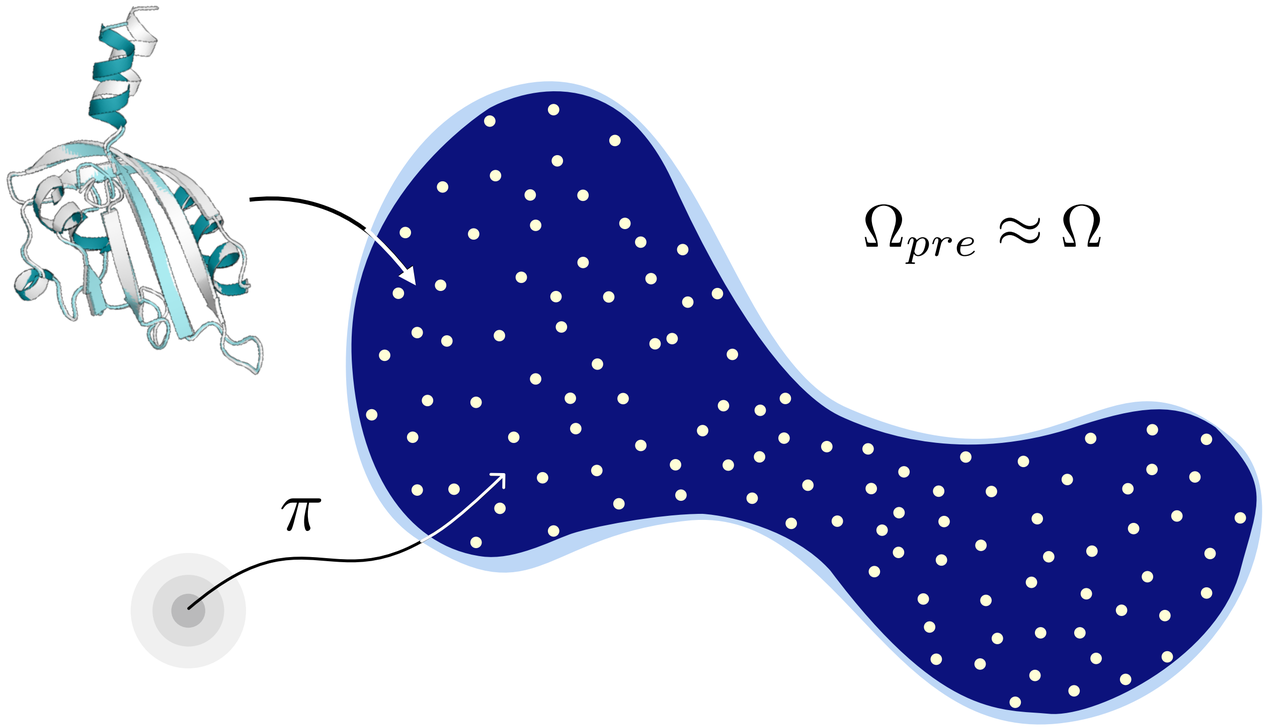


Data $D = \{x_i\}_{i=1}^{n}$

Generative Discovery Paradigm

valid design space Ω

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



$\Omega_{pre} \approx \Omega$

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

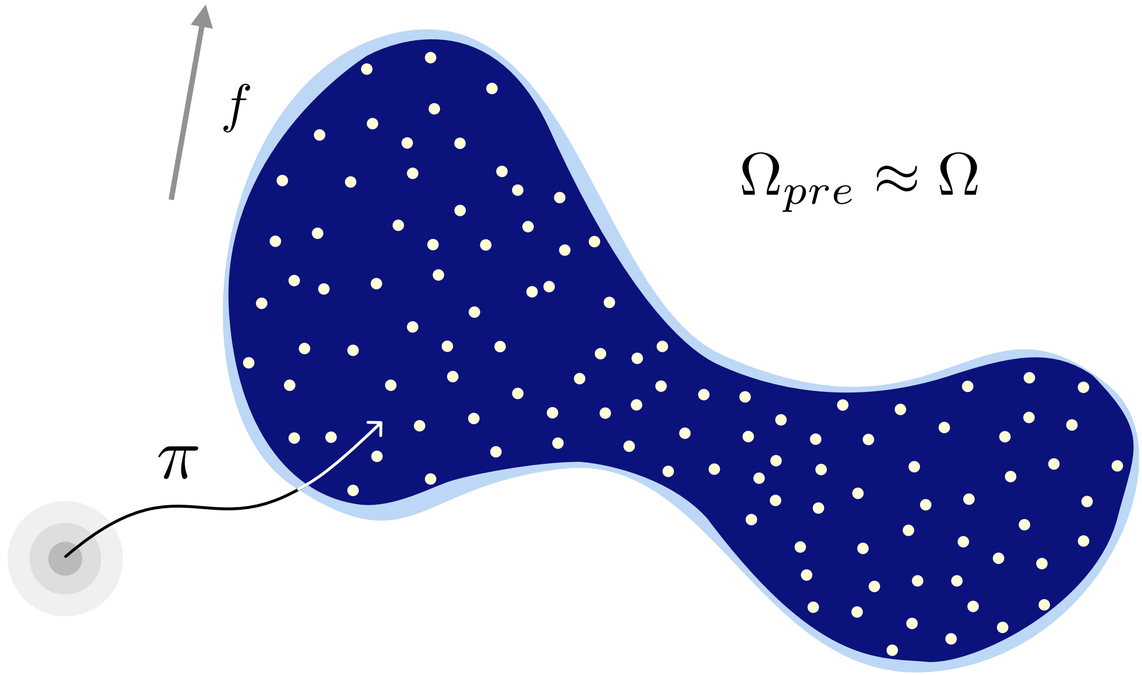
↓

Pre-trained generative model π

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

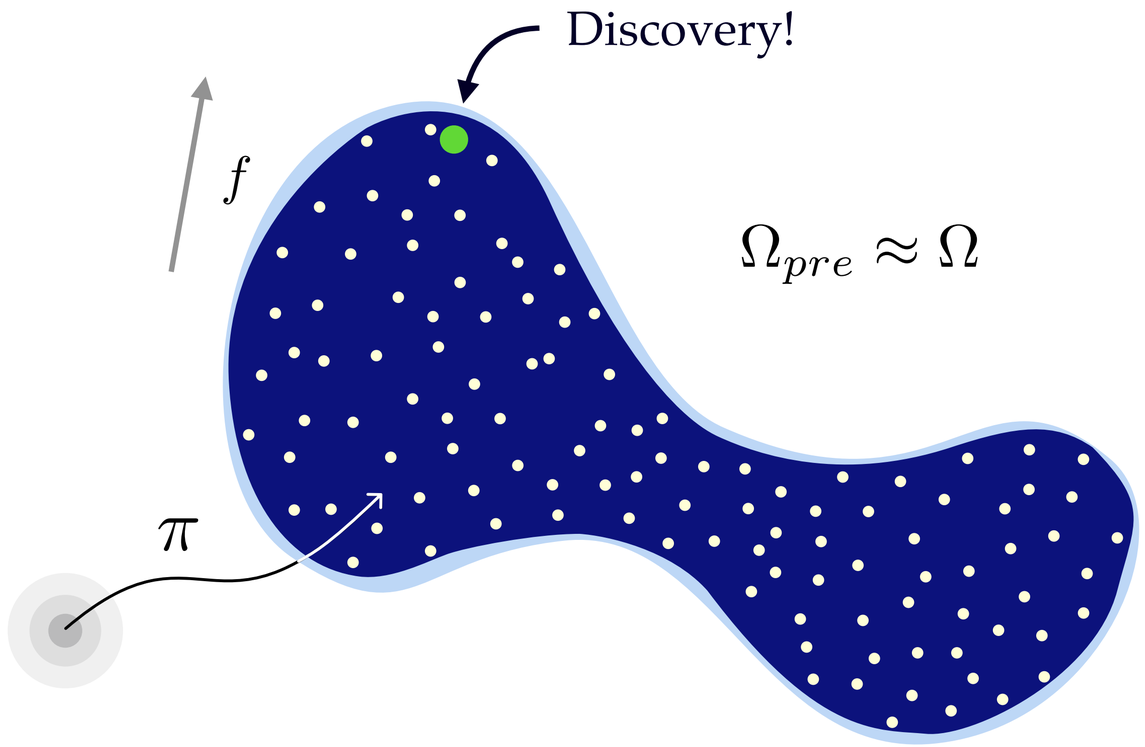
Discovery problem

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

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

Discovery problem

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

Generative Modeling

Learn generative prior from data

Generative Modeling

Learn generative prior from data



Current Status

Generative model in-distribution reward-adaptation to improve average behavior

Generative Modeling

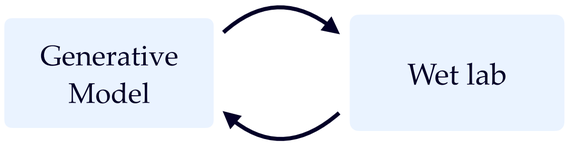
Learn generative prior from data

↓

Current Status

Generative model in-distribution reward-adaptation to improve average behavior

'discover novel peptides with promising anticancer activity'



Generative Modeling

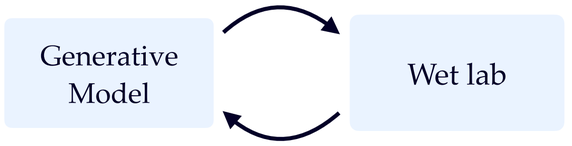
Learn generative prior from data

↓

Current Status

Generative model in-distribution reward-adaptation to improve average behavior

'discover novel peptides with promising anticancer activity'



not about in-distribution samples or average behavior and with limited budget of experiments

Generative Modeling

Learn generative prior from data

↓

Current Status

Generative model in-distribution reward-adaptation to improve average behavior

↓ This talk

Generative Discovery

Generation of **rare and diverse structures,** potentially **far beyond available data** and **sample-efficiently**

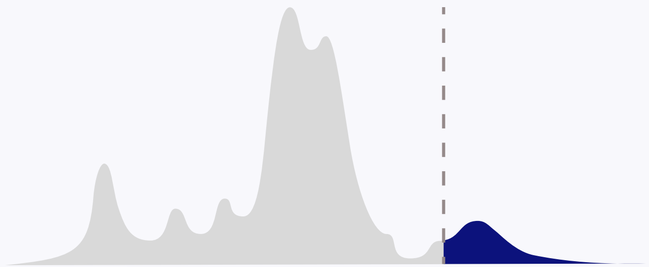
Generative Discovery

Generation of **rare and diverse structures,** potentially **far beyond available data** and **sample-efficiently**

Part I

Rare and Diverse Generation

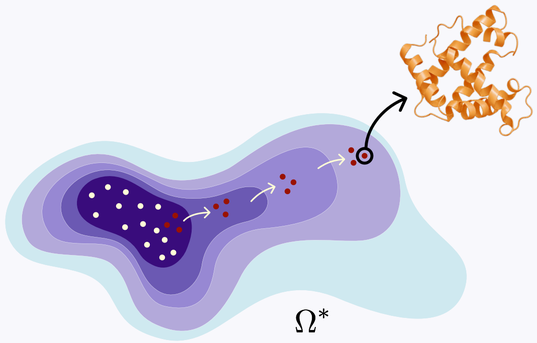
Which candidates matter?



Part II

Out-of-Distribution Flow Modeling

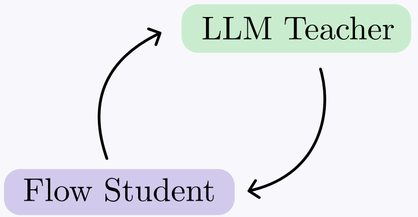
Which regions can we reach?



Part III

Sample-Efficient Discovery via LLM Feedback

What feedback guides us there?



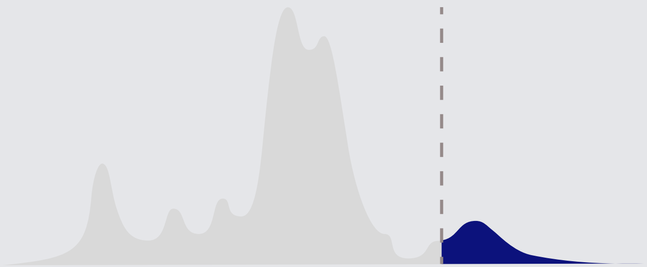
This talk:

Generative Discovery Beyond the Data

Part I

Rare and Diverse Generation

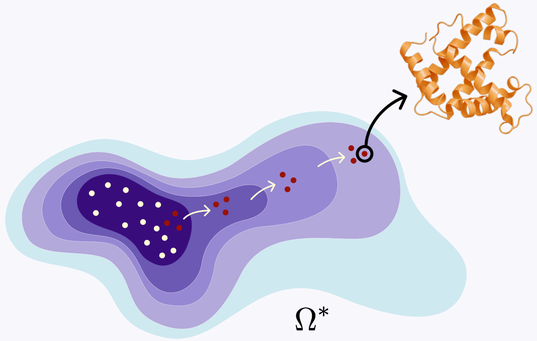
Which candidates matter?



Part II

Out-of-Distribution Flow Modeling

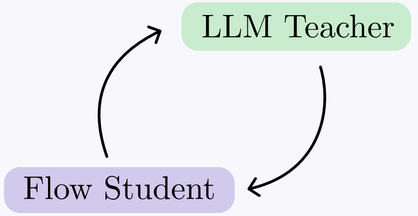
Which regions can we reach?



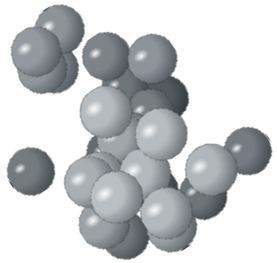
Part III

Sample-Efficient Discovery via LLM Feedback

What feedback guides us there?



Diffusion and Flow Modeling



noise
distribution

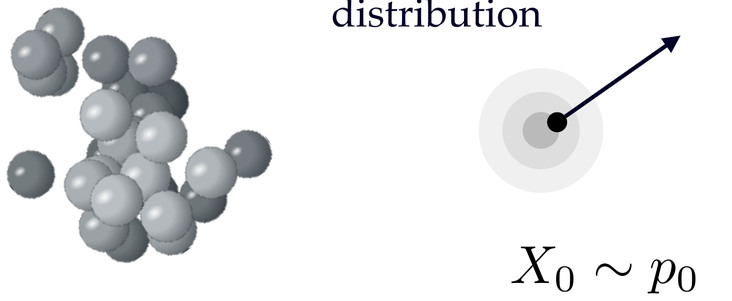


$X_0 \sim p_0$

Generation via learned vector field

Sample noise $X_0 \sim p_0$ $\longrightarrow$ Follow learned vector field u^{pre}

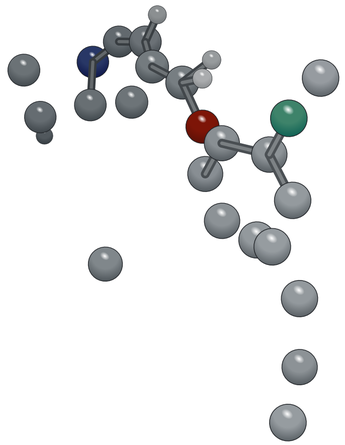
Diffusion and Flow Modeling

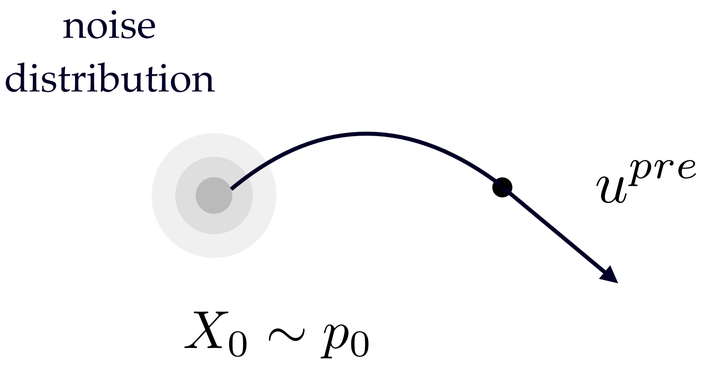


Generation via learned vector field

Sample noise $X_0 \sim p_0$ $\longrightarrow$ Follow learned vector field u^{pre}

Diffusion and Flow Modeling

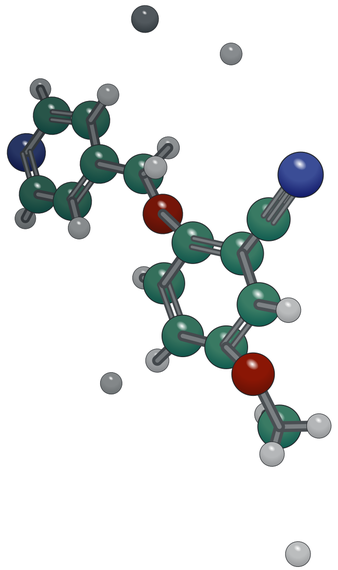


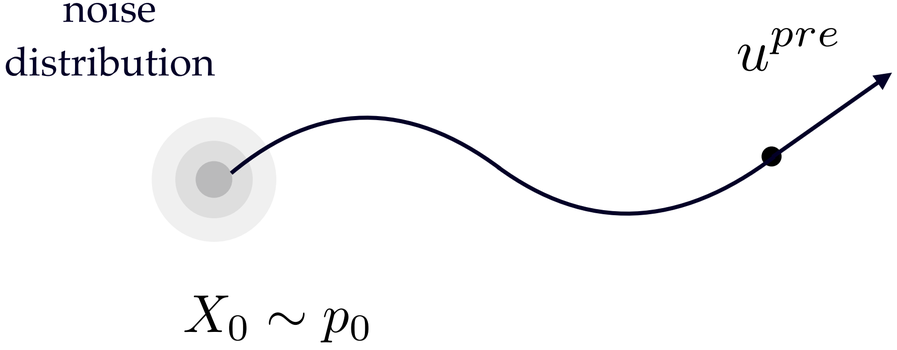


Generation via learned vector field

Sample noise $X_0 \sim p_0$ ⟶ Follow learned vector field u^{pre}

Diffusion and Flow Modeling

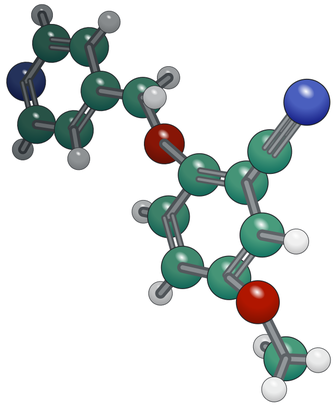


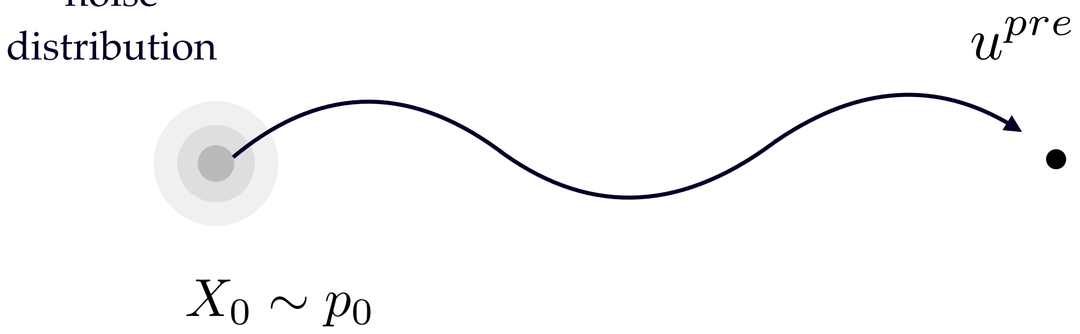


Generation via learned vector field

Sample noise $X_0 \sim p_0$ $\longrightarrow$ Follow learned vector field u^{pre}

Diffusion and Flow Modeling

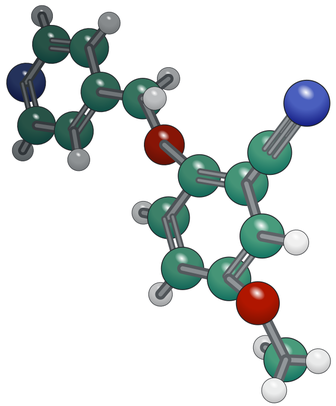




Generation via learned vector field

Sample noise $X_0 \sim p_0$ $\longrightarrow$ Follow learned vector field u^{pre}

Diffusion and Flow Modeling



noise distribution

$X_0 \sim p_0$

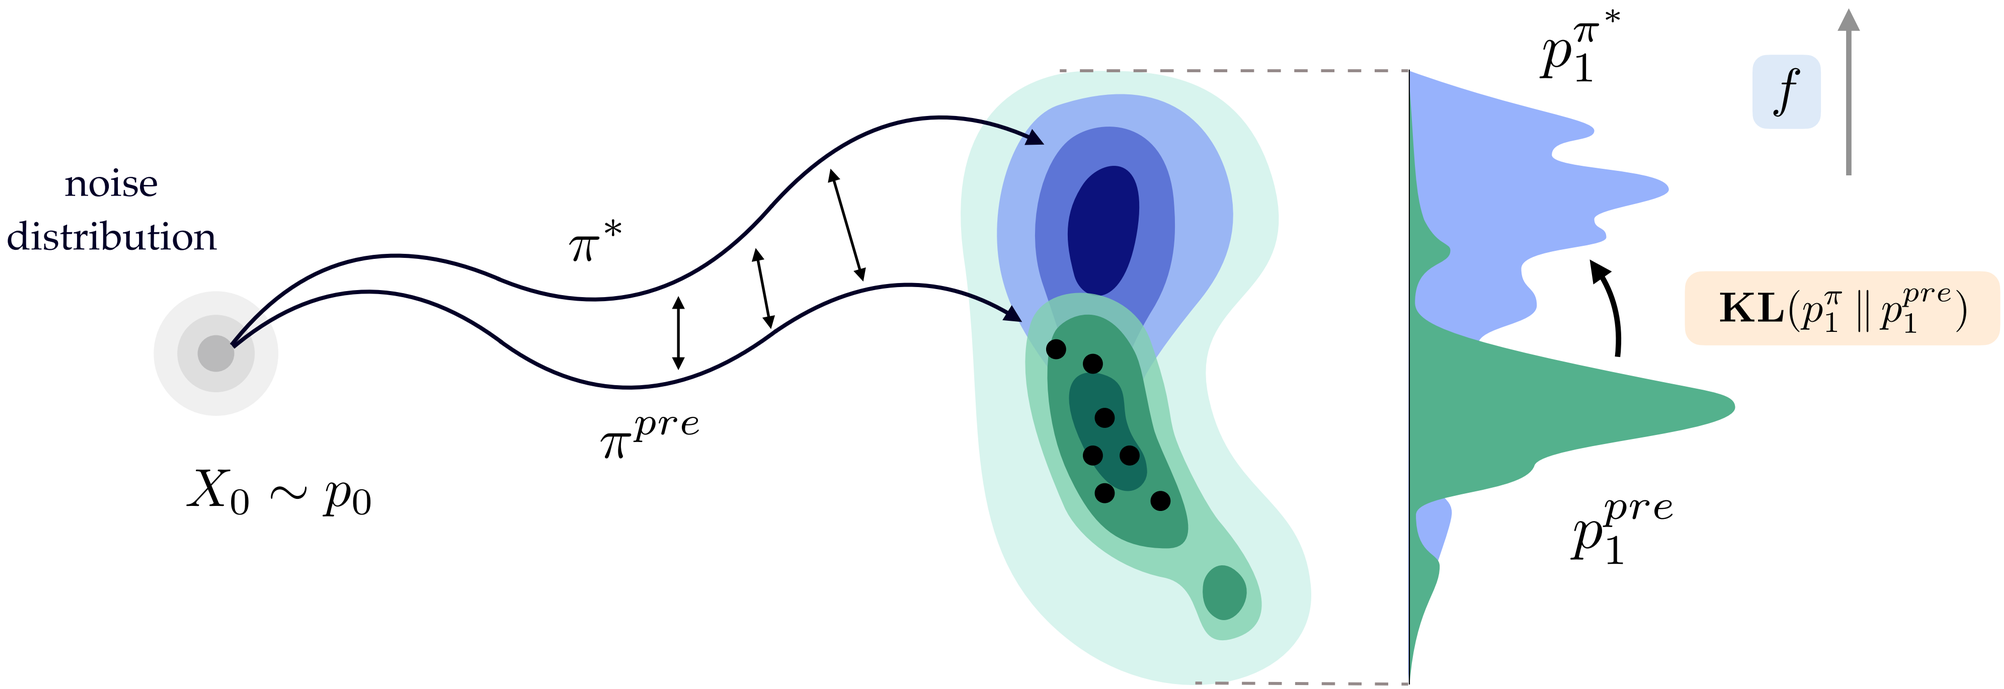
u^{pre}

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

Generation via learned vector field

Sample noise $X_0 \sim p_0$ ⟶ Follow learned vector field u^{pre}

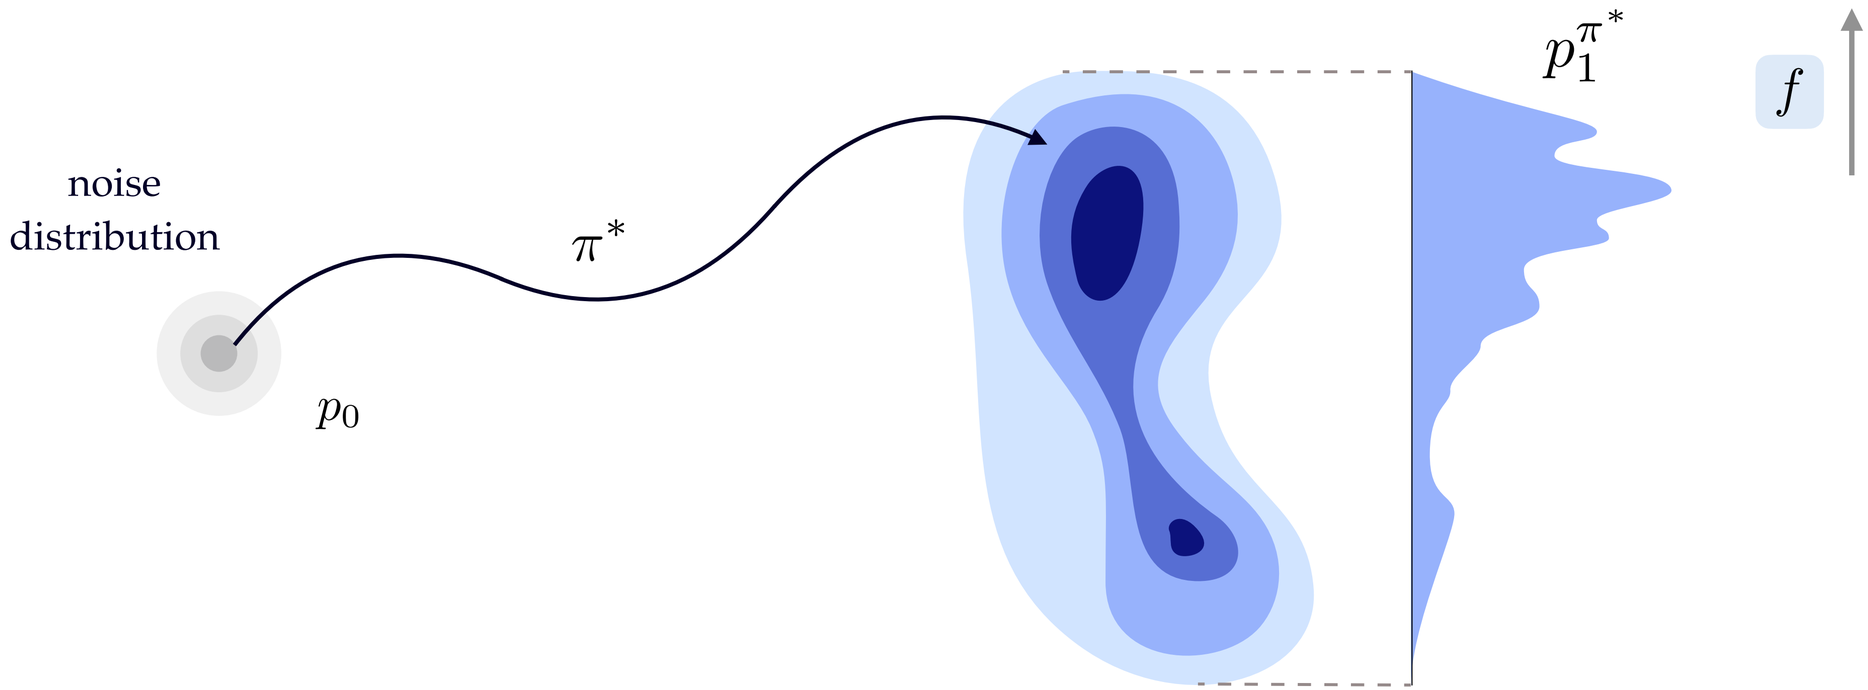
Generative Optimization via Flow Adaptation



Flow Reward-Guided Adaptation

$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \quad \mathbb{E}_{x\sim p_1^\pi}\left[f(x)\right] - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

Generative Optimization via Flow Adaptation



Flow Reward-Guided Adaptation

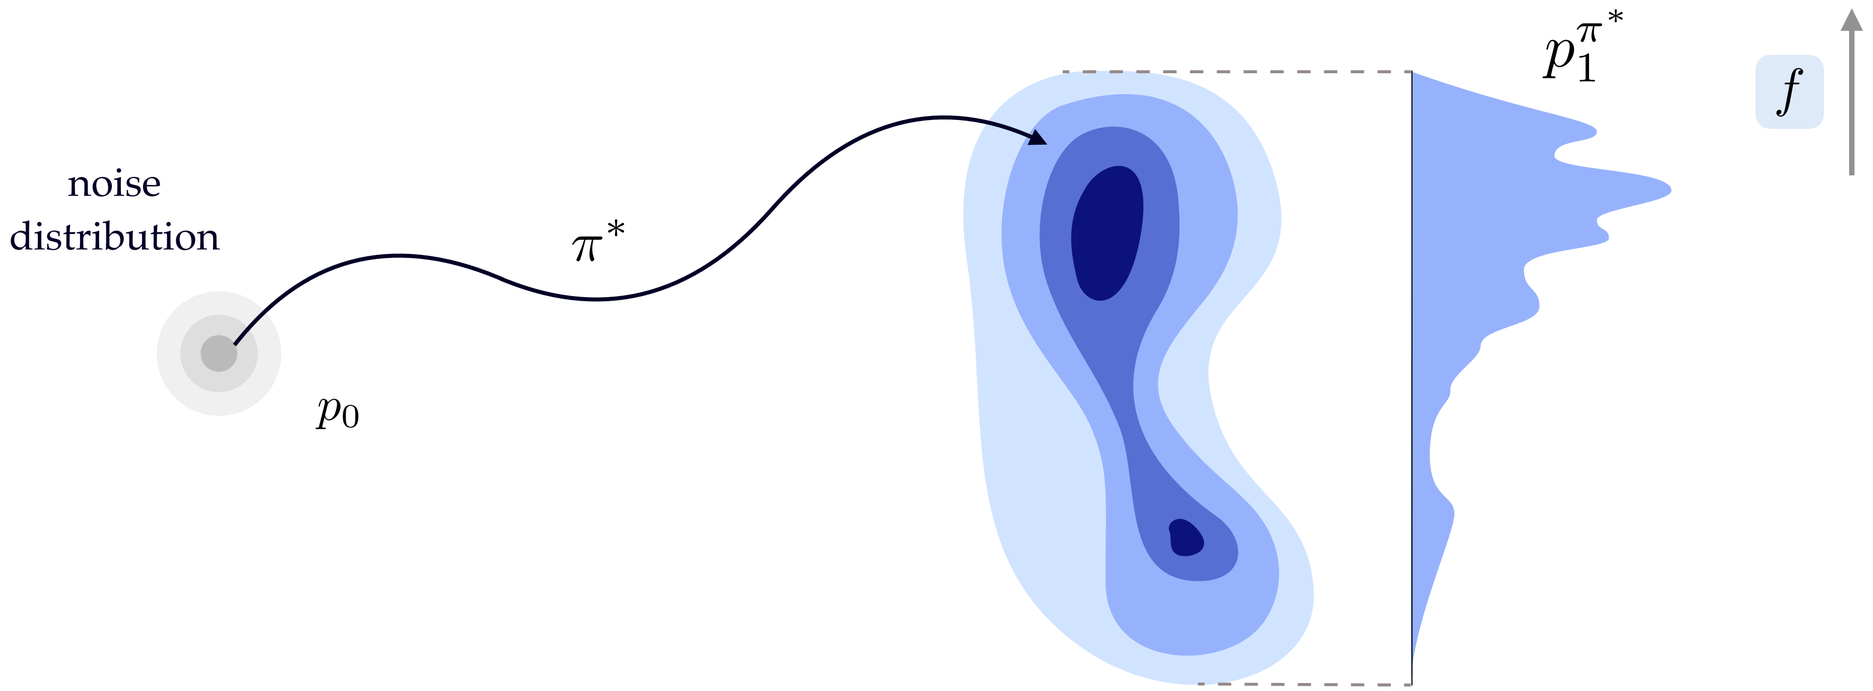
$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \quad \mathbb{E}_{x\sim p_1^\pi}\left[f(x)\right] - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

Closed form *reward tilted distribution*

$$p_1^{\pi^*}(x) \propto \exp\left(\frac{f(x)}{\alpha}\right) \cdot p_1^{pre}(x)$$

in-distribution density reweighing

Generative Optimization via Flow Adaptation



Flow Reward-Guided Adaptation

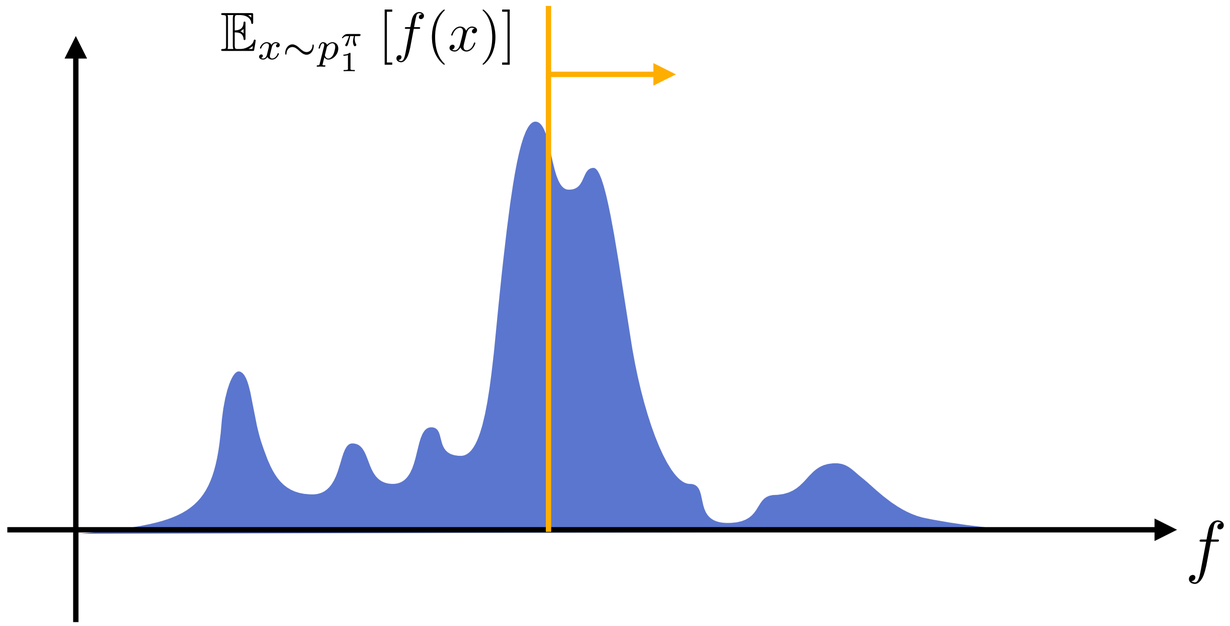
$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi}[f(x)] - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

E.g.,

FlowGRPO (RL, **gradient-free**)

Adjoint Matching (Control, **gradient-based**)

Is Expended Reward Maximization Good for Discovery?

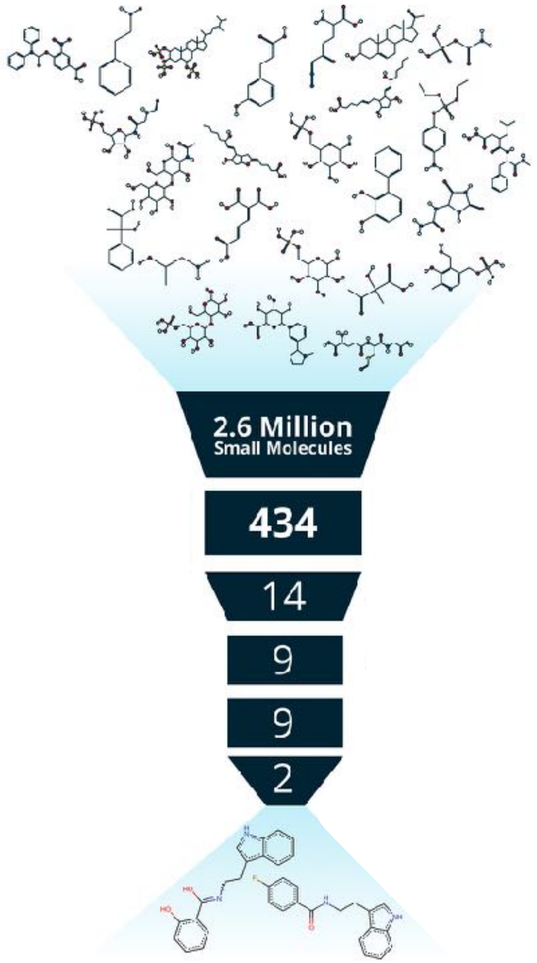


Flow Reward-Guided Fine-Tuning

$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \; \mathbb{E}_{x\sim p_1^\pi}[f(x)] - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

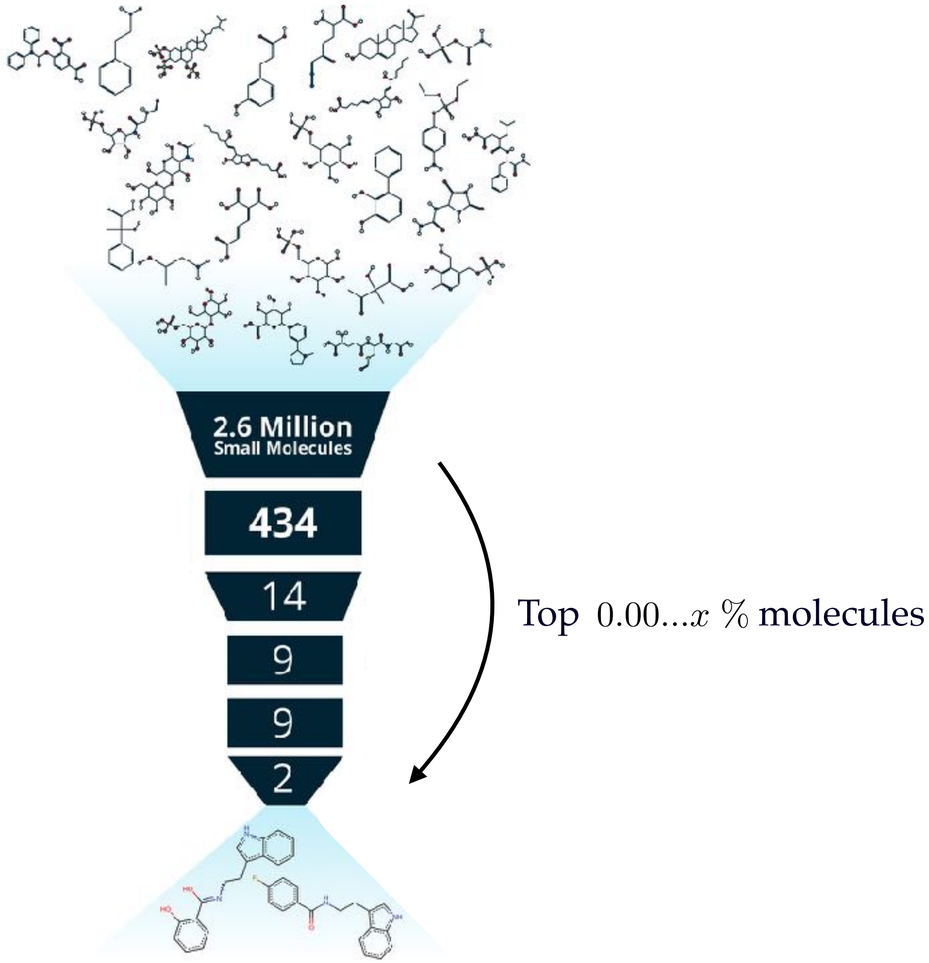
Optimizes average sample quality

Discovery Pipeline in Chemistry and Biology



[Anastasiia Gryniukova et al., 2023]

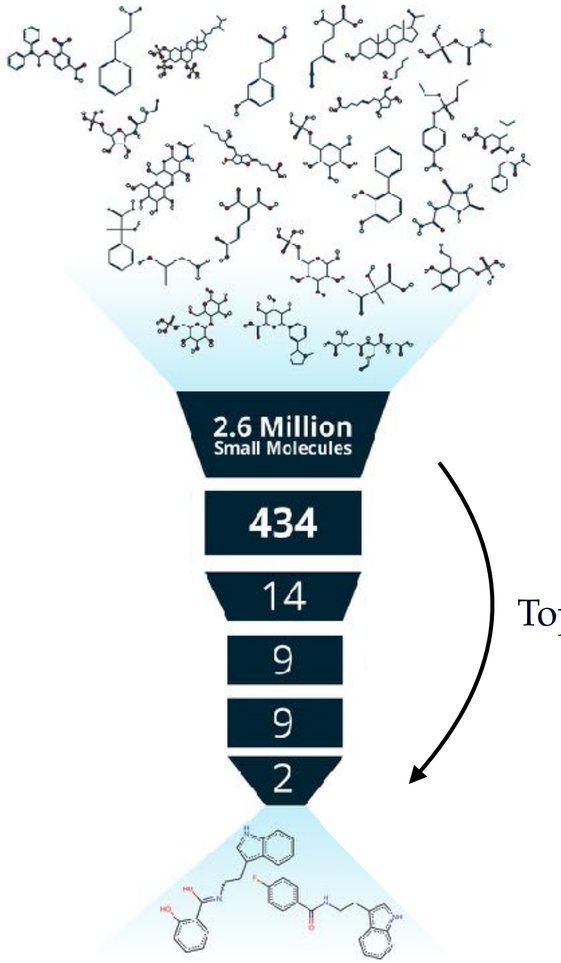
Discovery Pipeline in Chemistry and Biology



Top $0.00...x$ % molecules

[Anastasiia Gryniukova et al., 2023]

Discovery Pipeline in Chemistry and Biology



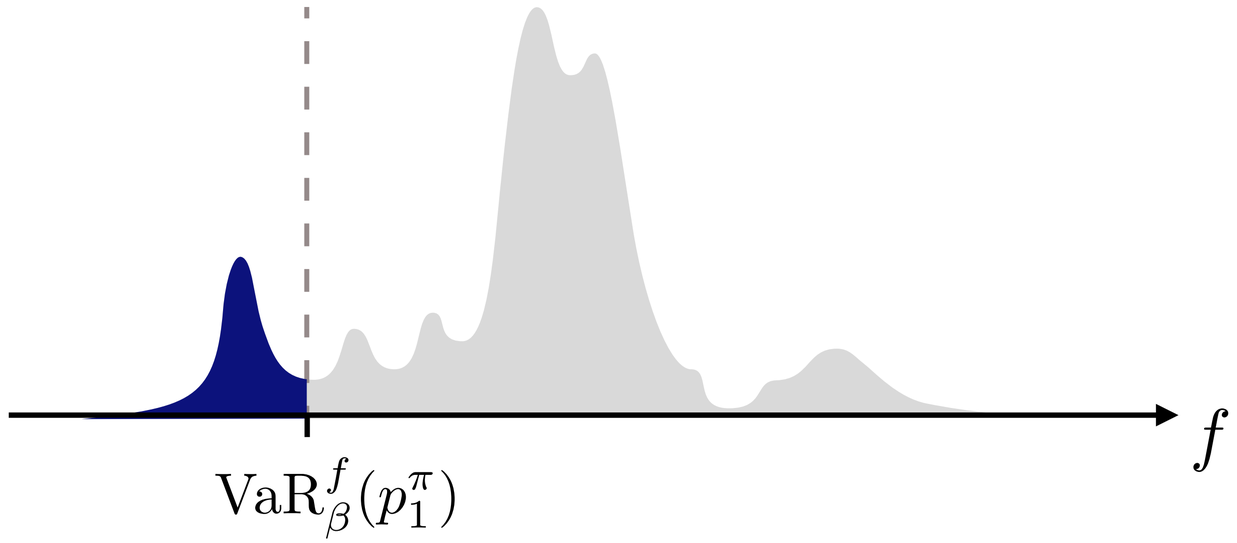
Top $0.00...x$ % molecules

Takeaway

Average reward adaptation is often a bad objective for discovery.

[Anastasiia Gryniukova et al., 2023]

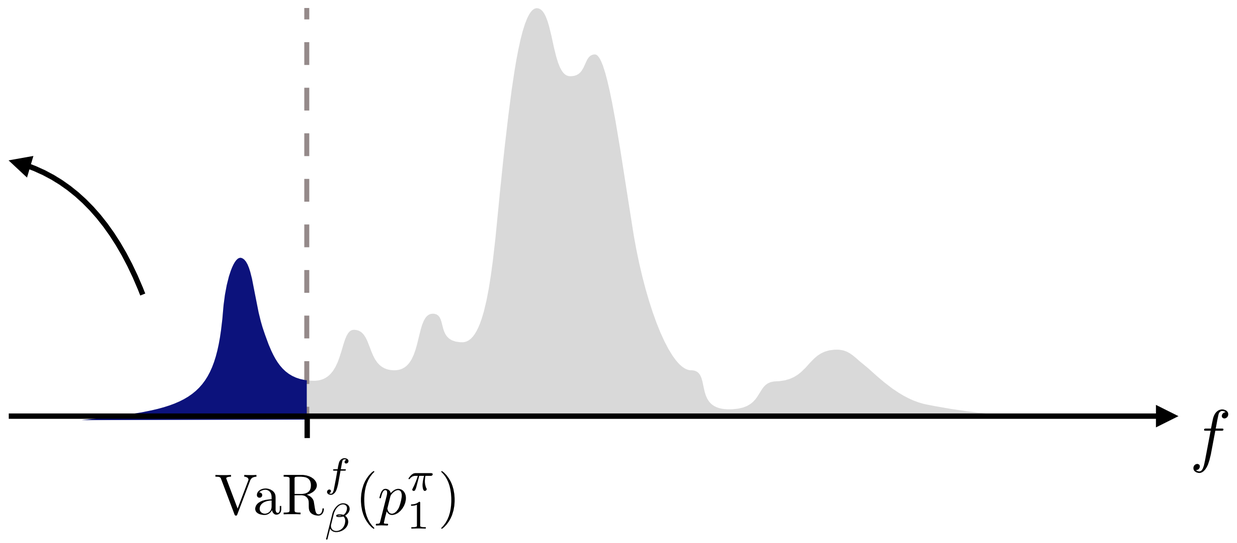
Risk-sensitivity for Quantitative Finance



Risk-sensitivity for Quantitative Finance



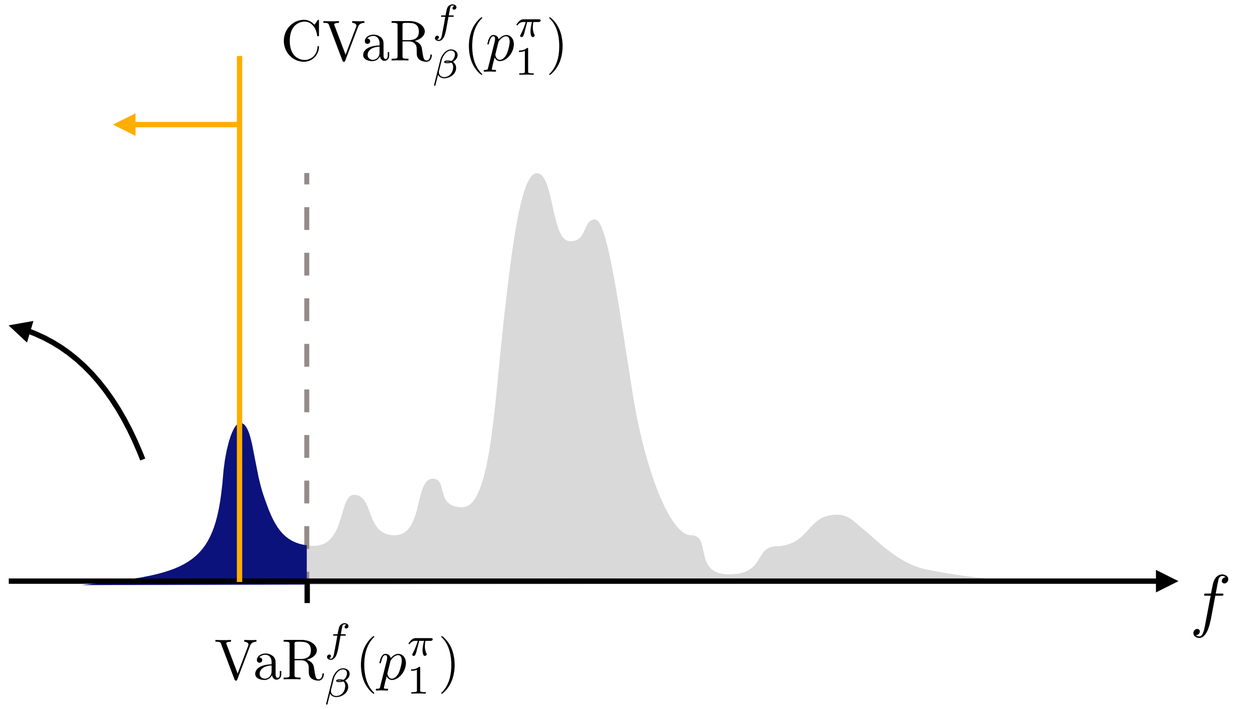




Risk-Sensitivity for Quantitative Finance



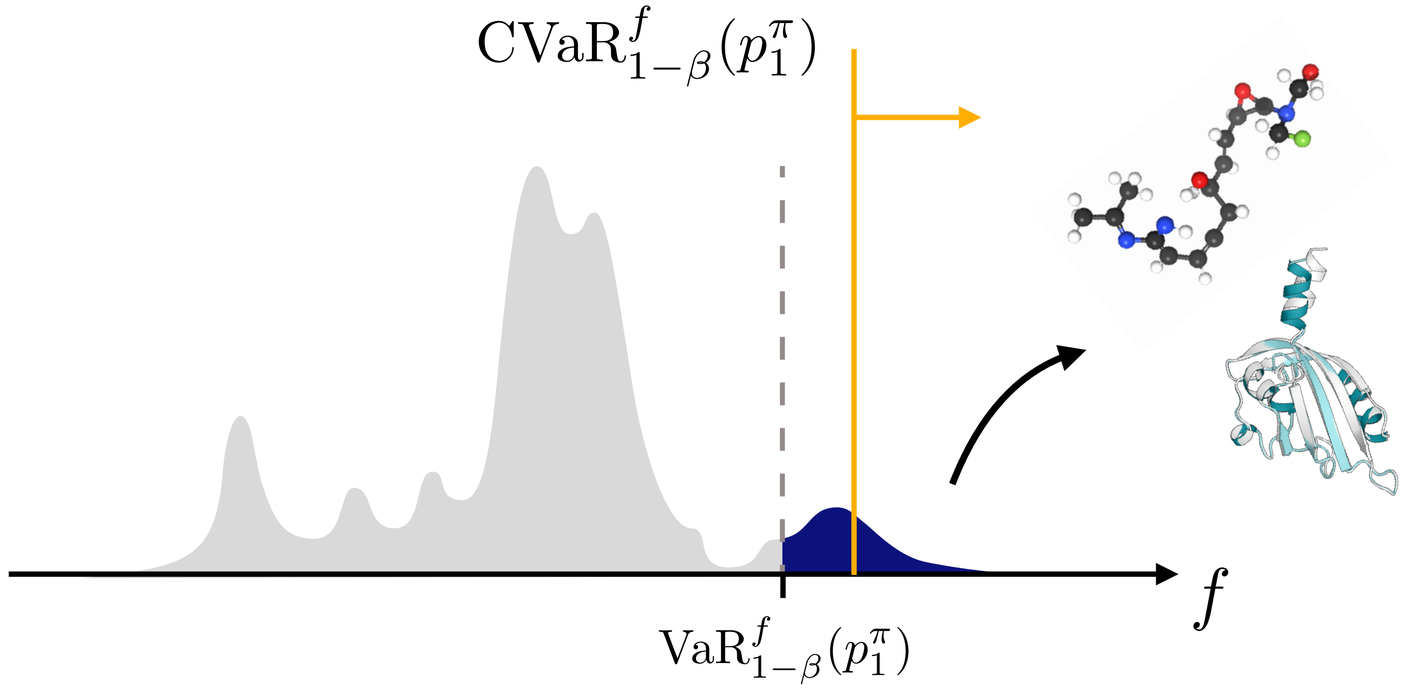




$$\mathrm{CVaR}^f_\beta(p^\pi_1) := \mathbb{E}_{x\sim p^\pi_1}[f(x) \mid f(x) \leq \mathrm{VaR}^f_\beta(p^\pi_1)]$$

Expected value of **left (bottom)** β **quantile**

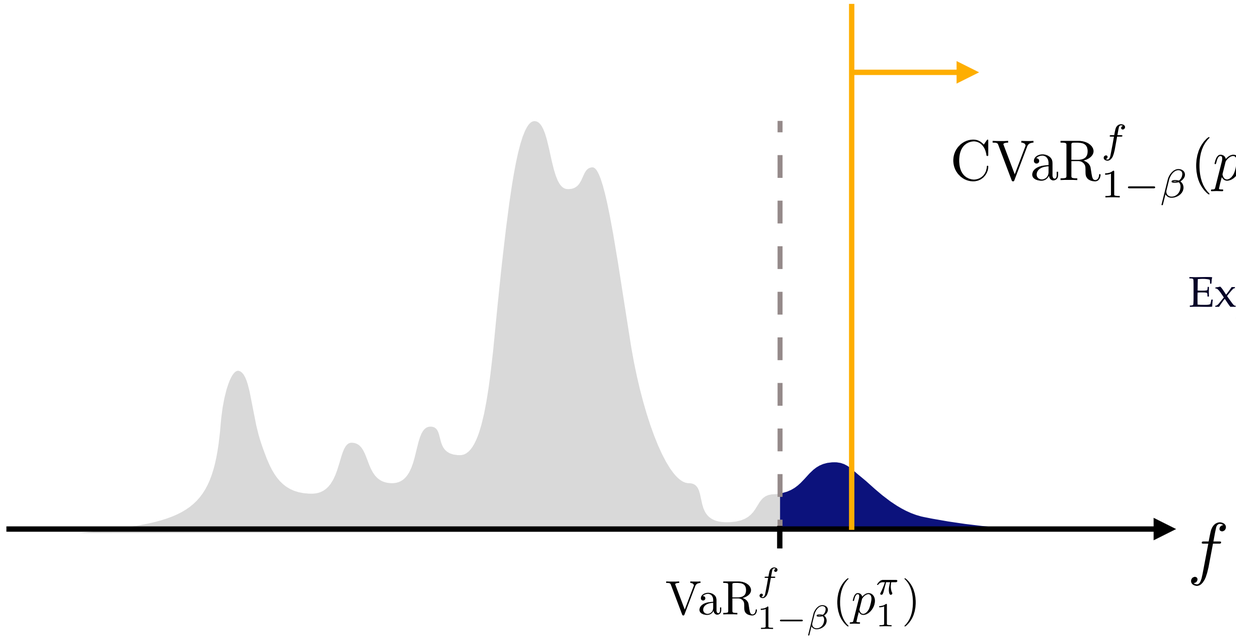
Risk-Sensitivity for Scientific Discovery



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

Expected value of **right (top)** $1-\beta$ **quantile**

Tail-Aware Flow Adaptation for Discovery



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

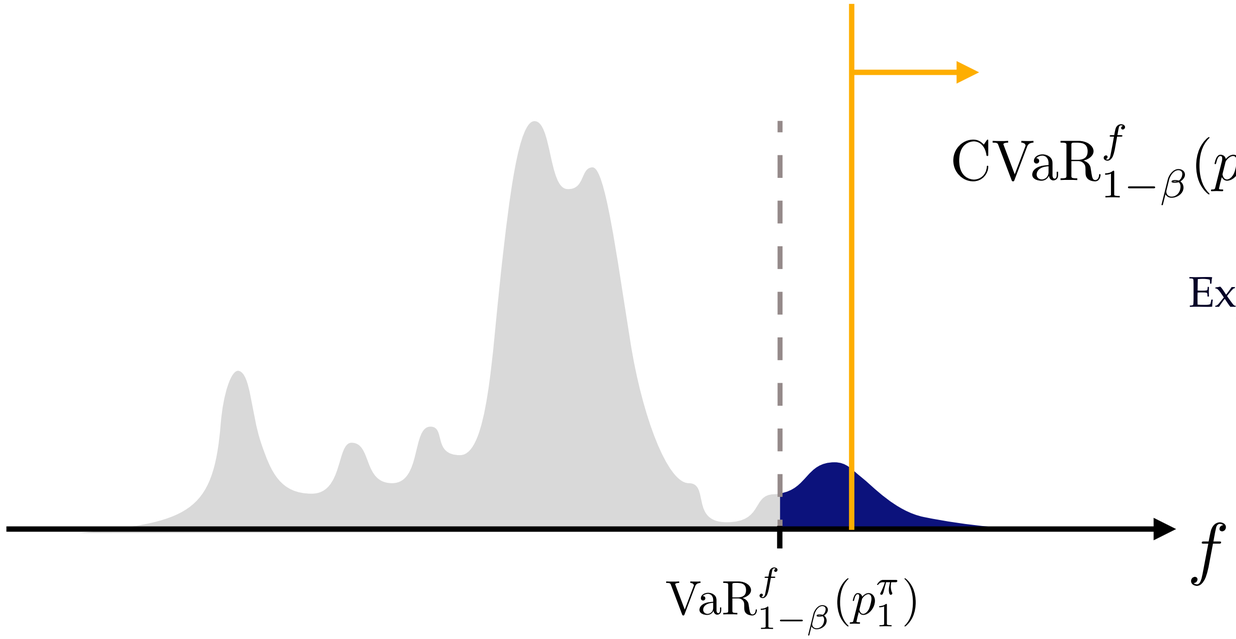
Expected value of top $1-\beta$ quantile

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \operatorname*{argmax}_{\pi} \mathrm{CVaR}_{1-\beta}(p_1^{\pi}) - \alpha \mathbf{KL}(p_1^{\pi} \,\|\, p_1^{pre})$$

→ **Prior regularization to generate valid structures**

Tail-Aware Flow Adaptation for Discovery



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

Expected value of top $1-\beta$ quantile

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \underset{\pi}{\mathrm{argmax}} \; \mathrm{CVaR}_{1-\beta}(p_1^{\pi}) - \alpha \mathbf{KL}(p_1^{\pi} \,\|\, p_1^{pre})$$

→ **Prior regularization to generate valid structures**

CVaR is a **non-linear** **functional. How to optimize?**

CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \quad \mathrm{CVaR}_{1-\beta}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \operatorname*{argmax}_{\pi} \quad \mathrm{CVaR}_{1-\beta}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

$$\mathrm{CVaR}_{1-\beta}(p_1^\pi) = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{\beta} \mathbb{E}_{x \sim p_1^\pi} [(f(x) - t)_+] \right\}$$

CVaR variational formulation

CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \operatorname*{argmax}_{\pi} \quad \mathrm{CVaR}_{1-\beta}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

$$\mathrm{CVaR}_{1-\beta}(p_1^\pi) = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{\beta} \mathbb{E}_{x \sim p_1^\pi} [(f(x) - t)_+] \right\}$$

CVaR variational formulation

First solve a scalar threshold problem:

$$t^* = \arg\min_{t \in \mathbb{R}} \left\{ t + \alpha \log \mathbb{E}_{x \sim p_1^{pre}} \left[\exp\left(\frac{(f(x) - t)_+}{\alpha \beta} \right) \right] \right\}$$

1-dim opt. problem
(e.g., GD / line-search)

CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \operatorname*{argmax}_{\pi} \quad \mathrm{CVaR}_{1-\beta}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

$$\mathrm{CVaR}_{1-\beta}(p_1^\pi) = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{\beta} \mathbb{E}_{x \sim p_1^\pi} [(f(x) - t)_+] \right\}$$

CVaR variational formulation

First solve a scalar threshold problem:

$$t^* = \arg\min_{t \in \mathbb{R}} \left\{ t + \alpha \log \mathbb{E}_{x \sim p_1^{pre}} \left[\exp\left(\frac{(f(x) - t)_+}{\alpha \beta} \right) \right] \right\}$$

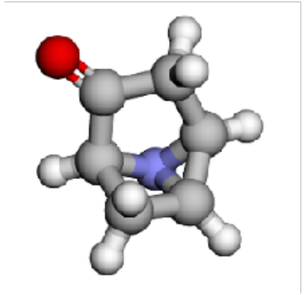
1-dim opt. problem
(e.g., GD / line-search)

Then fine-tune flow once:

$$r^*(x) = \frac{(f(x) - t^*)_+}{\beta} \longrightarrow \pi^* \in \operatorname*{argmax}_{\pi} \quad \mathbb{E}_{x \sim p_1^\pi} [r^*(x)] - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

~ 2 lines of code!

CVaR Flow Fine-Tuning for De Novo Molecular Design



GEOM-Drugs

$$g = (X, A, C, E)$$

Node Features
- Position (C)
- Atom Type (D)
- Formal Charge (D)

Edge Features
- Bond Type (D)

Flow over 3D molecular conformers

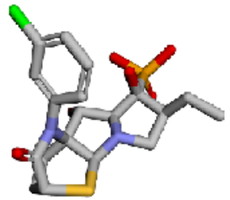
	Training time	FT calls	$\mathbb{E}[r]$	$\text{CVaR}_{0.9}[r]$	Validity	SA score
Pre-trained	0	0	$21.9_{\pm 1.21}$	$161.5_{\pm 10.6}$	$\mathbf{87.4\%}_{\pm 0.12\%}$	$7.64_{\pm 0.02}$
Adjoint M.	T	1	$\mathbf{29.8}_{\pm 0.58}$	$167.1_{\pm 7.22}$	$78.3\%_{\pm 1.01\%}$	$7.66_{\pm 0.02}$
TFFT	$\approx$ T	1	$27.1_{\pm 2.29}$	$\mathbf{183.4}_{\pm 15.8}$	$85.7\%_{\pm 0.59\%}$	$7.70_{\pm 0.04}$

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

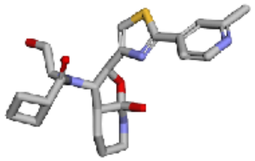
Reward = Negative energy via GFN1-xTB, Validity = RDKit Sanitizer

Top-Samples via CVaR Flow Fine-tuning: GEOM-Drugs

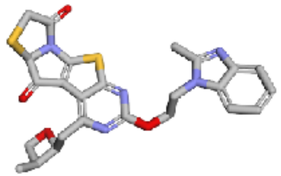
Pre-trained



840

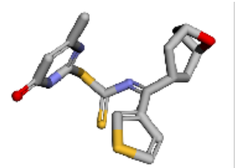


826

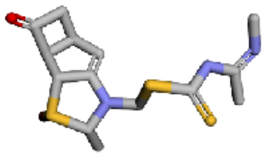


755

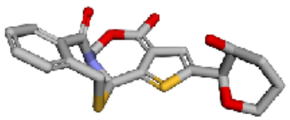
Adjoint Matching



1069



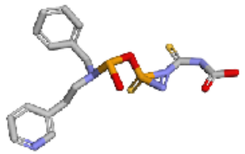
891



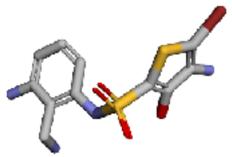
852

~1.3x higher-rewards

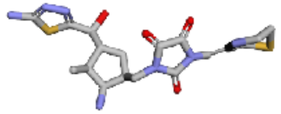
TFFT



5059



3237



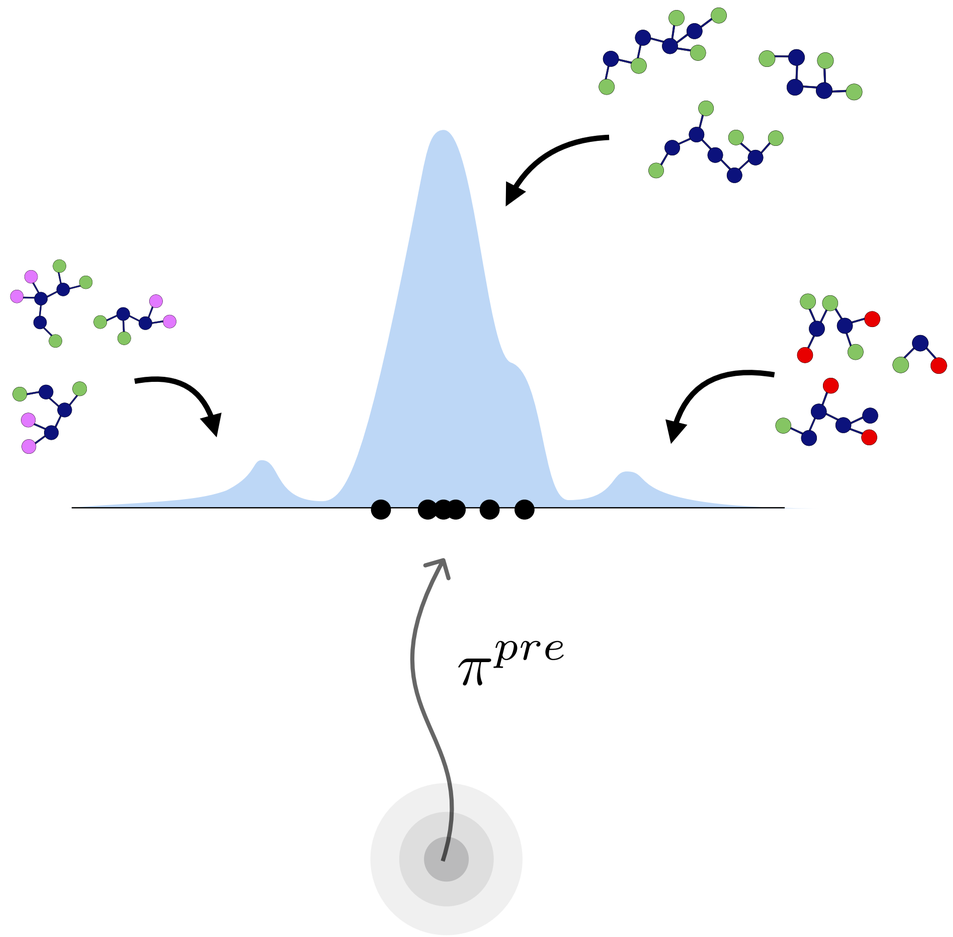
1032

~5x higher-rewards

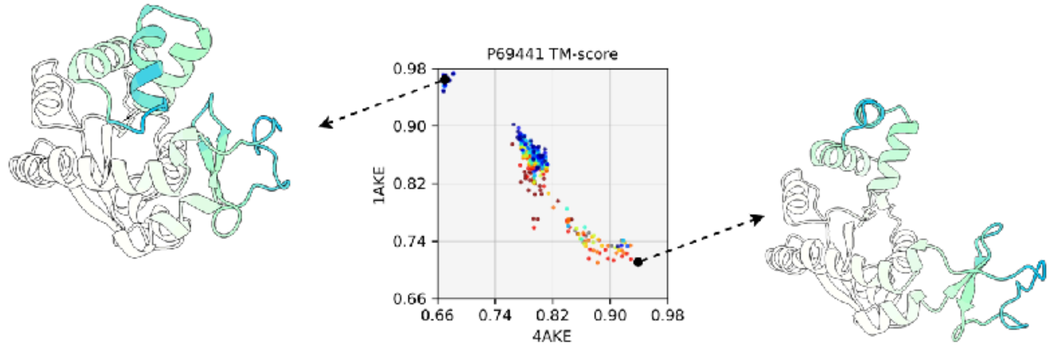
[Dunn and Koes, 2024] [Axelrod and Bombarelli, 2022]

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

Flow Adaptation to **Cover Diverse Modes**?

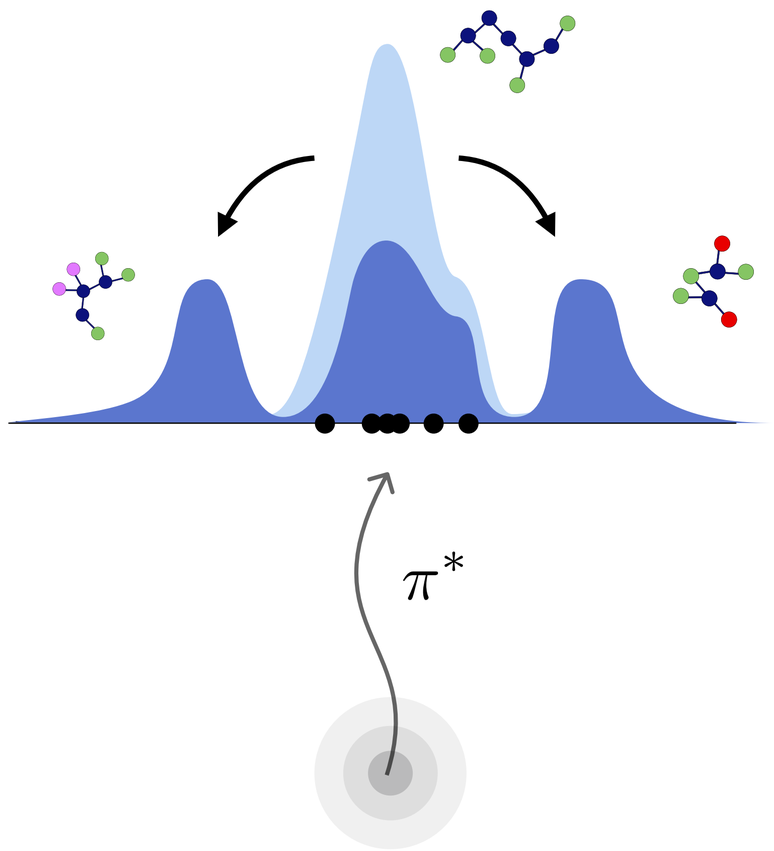


Key Hypothesis: protein generative models learn valid low-probability modes often hidden under finite samples.



[Shosuke Suzuki et al., 2026]

Flow Density Control for Diverse Modes Coverage



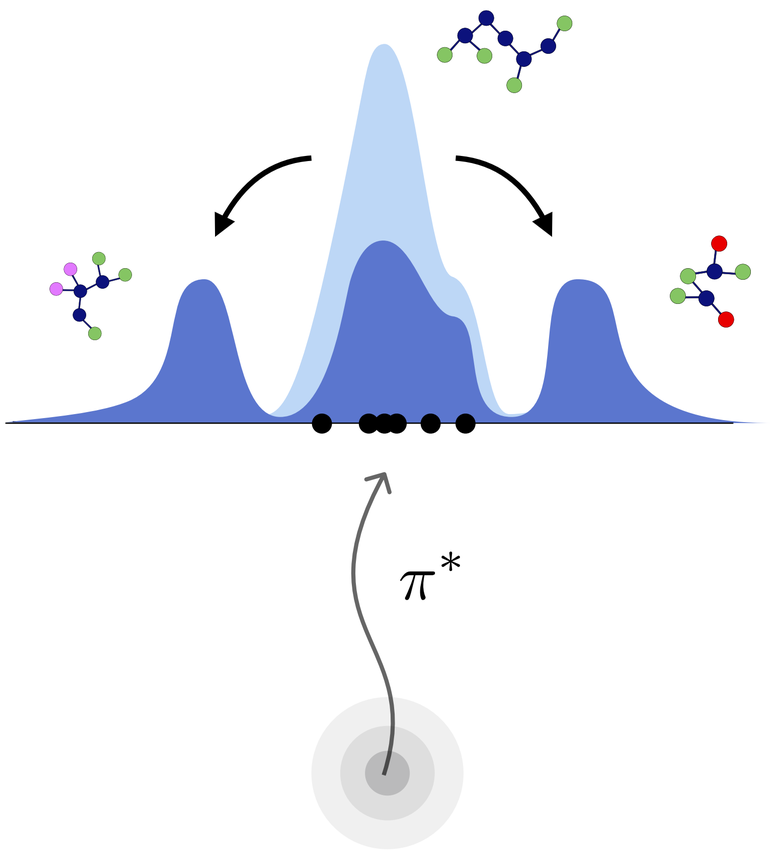
KL-regularized Entropic Expansion

$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \; \mathcal{H}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

Increases diversity

Avoid invalid regions

Flow Density Control for Diverse Modes Coverage



KL-regularized Entropic Expansion

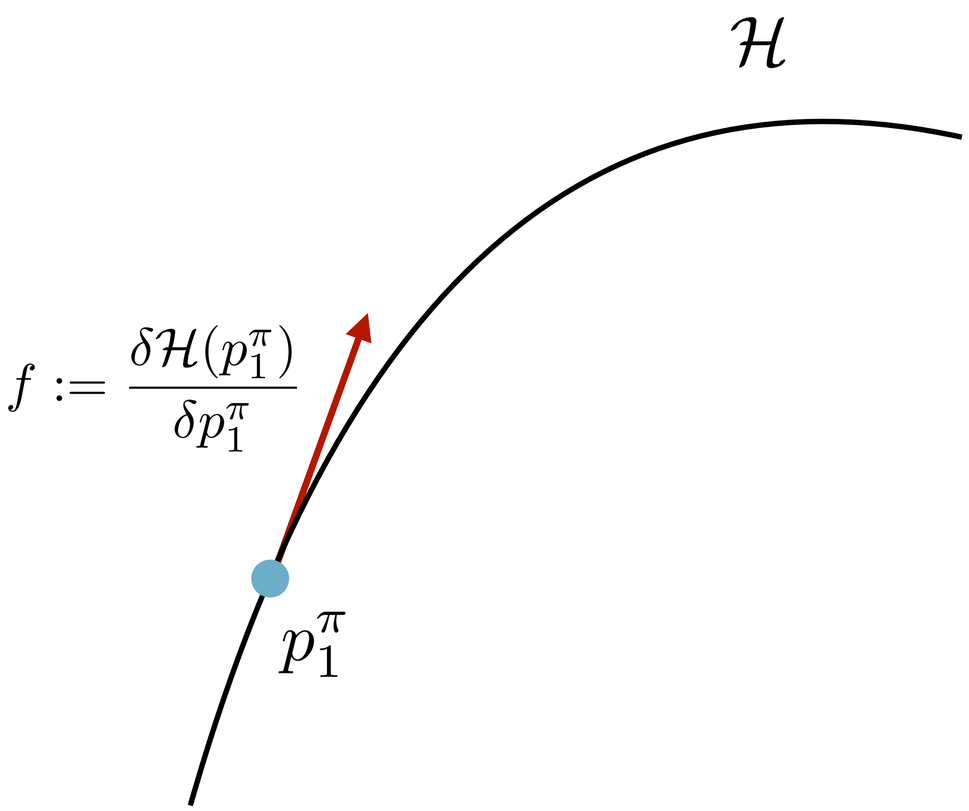
$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \; \mathcal{H}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$$

Increases diversity

Avoid invalid regions

Entropy is a **non-linear functional. How to optimize?**

Flow Density Control (FDC)

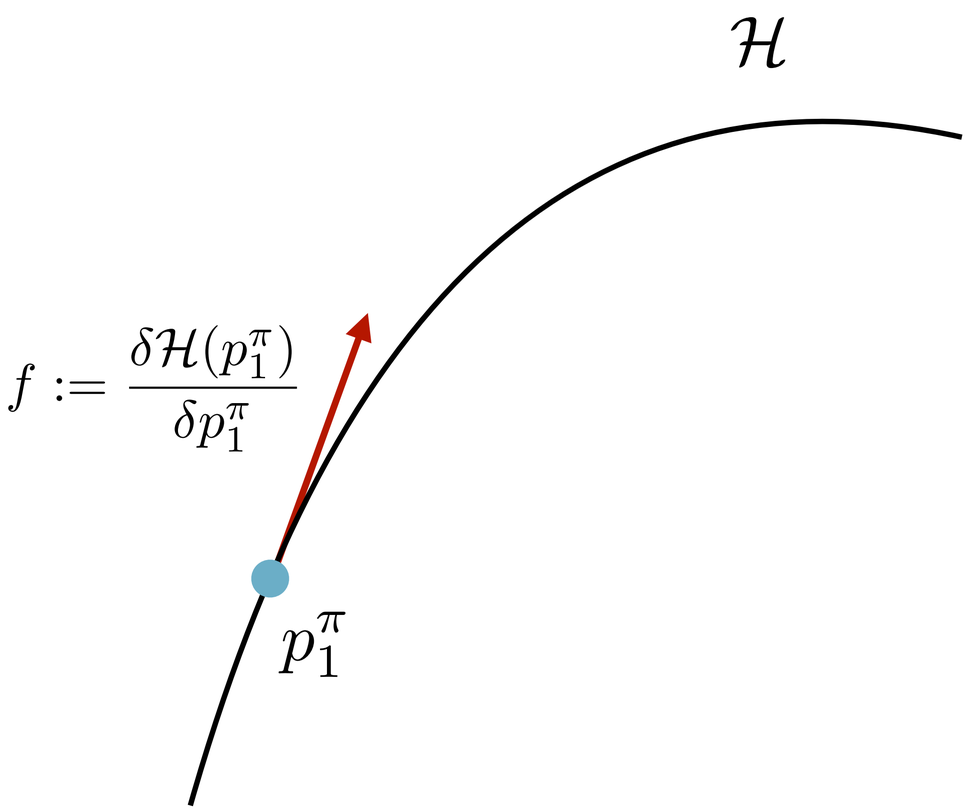


Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Flow Density Control (FDC)



Entropy First Variation (Surprise)

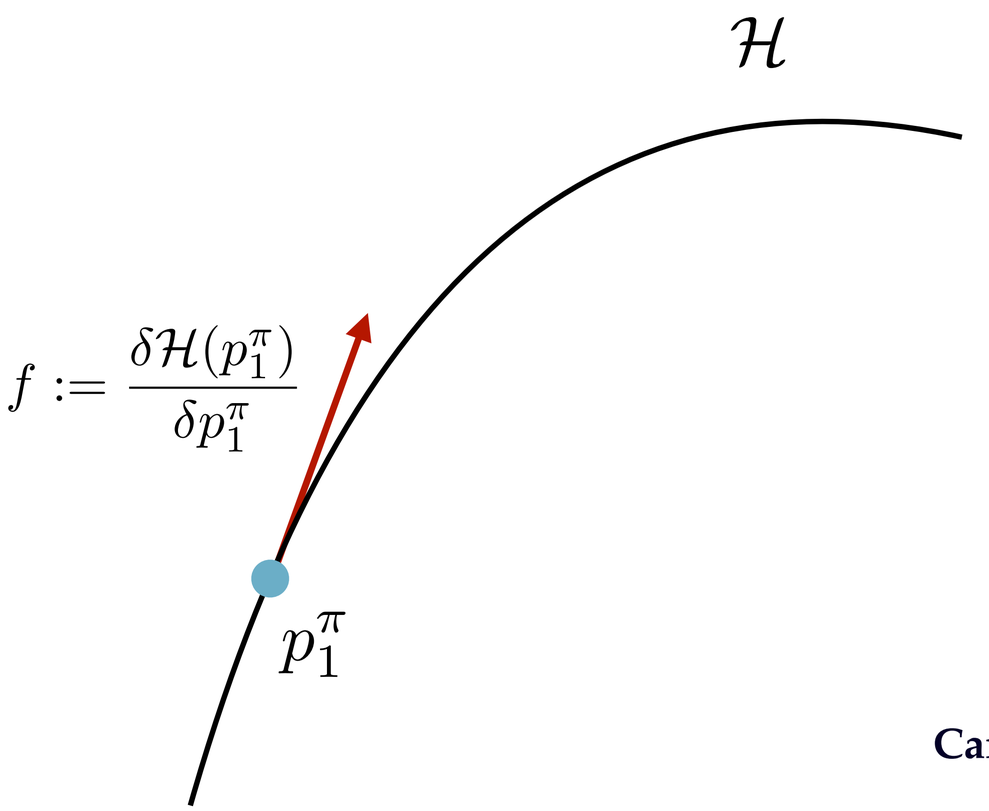
$$f := \frac{\delta \mathcal{H}\left(p_1^\pi\right)}{\delta p_1^\pi} = -\log\left(p_1^\pi\right)$$

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 Score

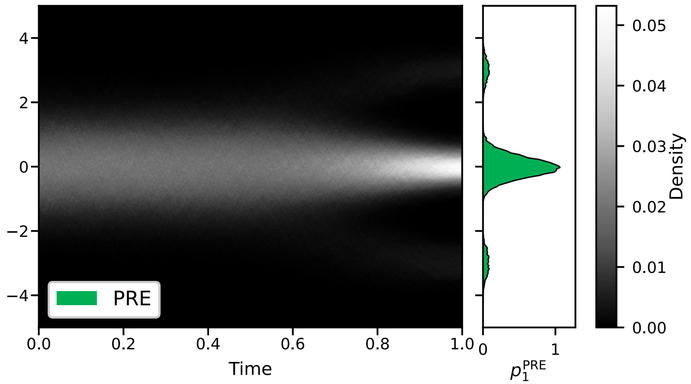
$$\nabla f = -\nabla \log\left(p_1^\pi\right) \approx -s^\pi(\cdot, 1)$$

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

Can use any gradient-based FT method.

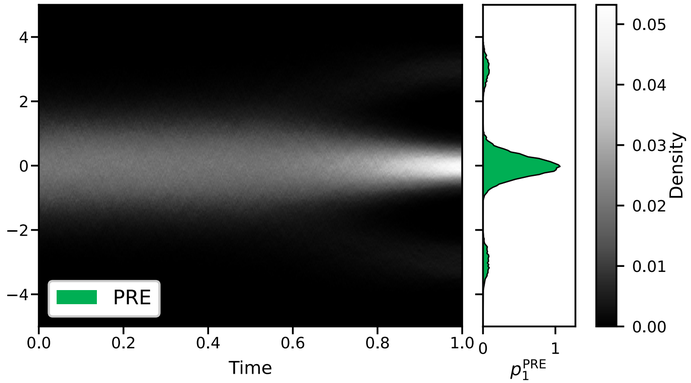
High-dimensional exploration with no estimation process!

Pre-trained generative model



$$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 1$$

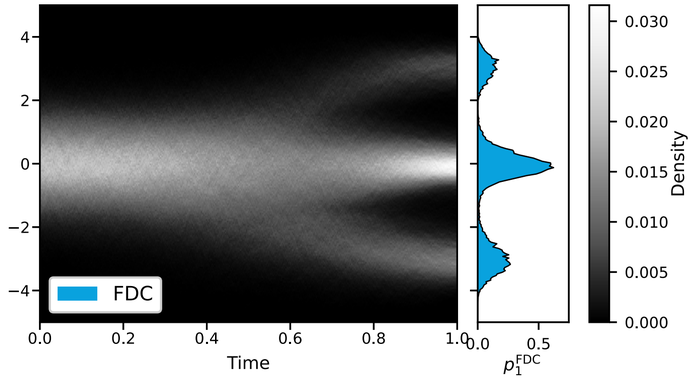
Pre-trained generative model



$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 1$

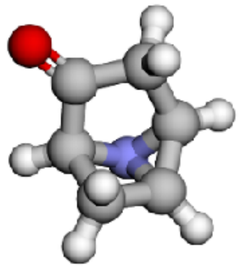
FDC

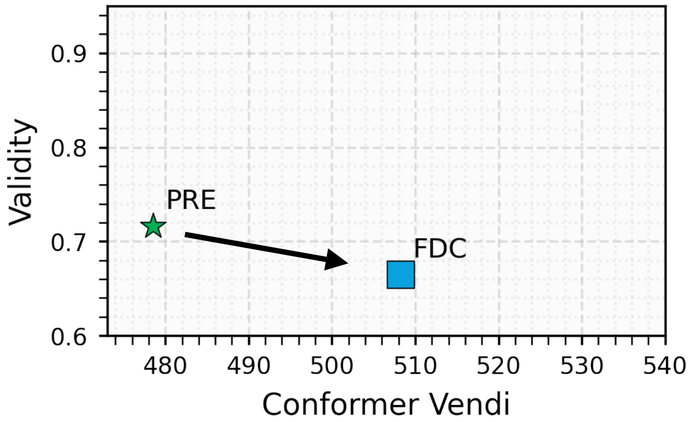
Expanded generative model



$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 0.5$

GEOM-Drugs

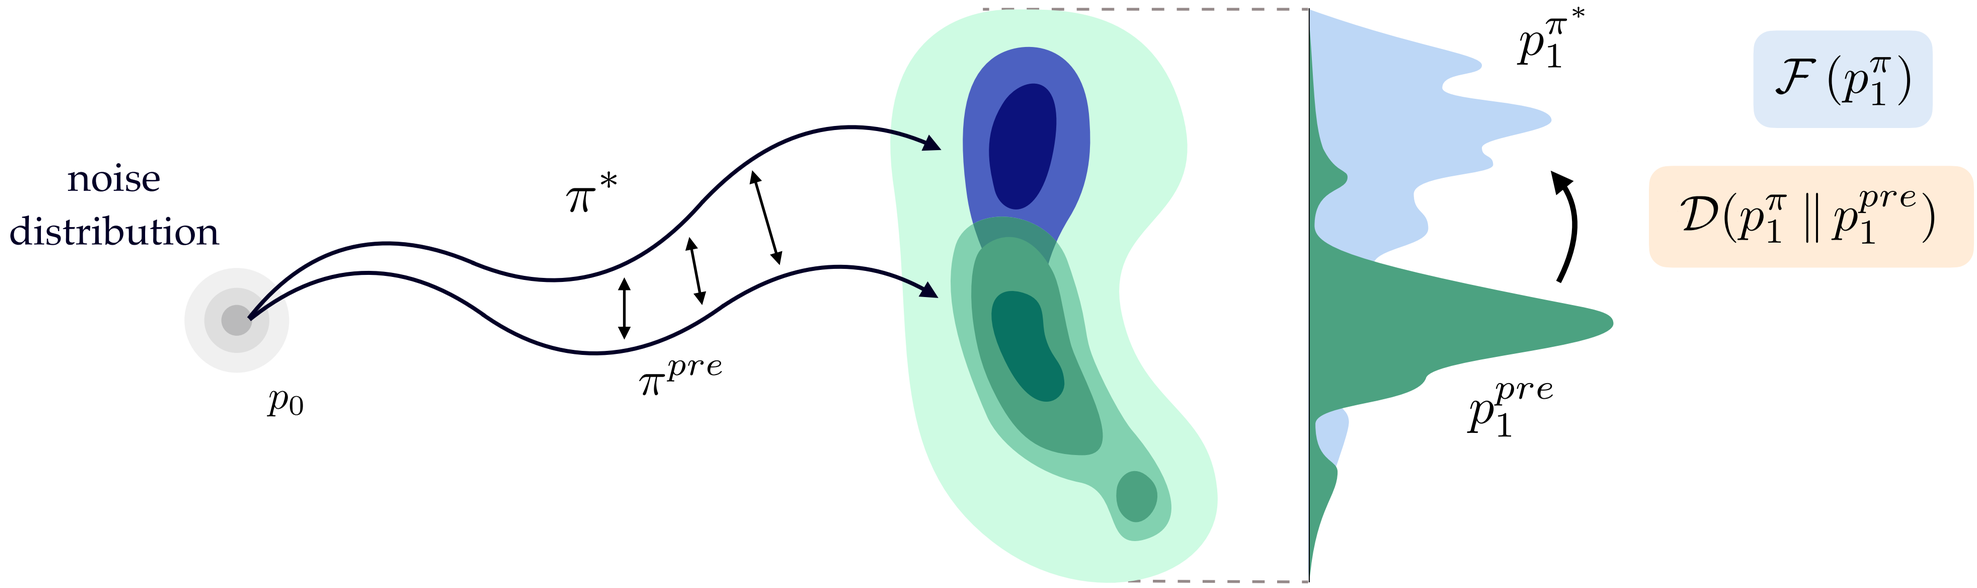




FDC achieves **higher conformer diversity**
of drug-like molecules

Flow Density Control for Distributional Fine-Tuning

Control of tail behavior is essential for discovery (e.g., diversity, tail-aware optimization)



How to represent new task? $\mathcal{F}$

How to preserve prior information? $\mathcal{D}$

Distributional Generative Optimization

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

Flow Density Control for Distributional Fine-Tuning

How to represent new task?

→ $\mathcal{F}$

How to preserve prior information?

→ $\mathcal{D}$

Application	Functional $\mathcal{F}$ / $\mathcal{D}$	Linear GO	Non-Linear GO	
			Convex	General
Reward optimization [14, 55]	$\mathbb{E}_{x\sim p^\pi}[r(x)]$	✓	✓	✓
Manifold Exploration [12] Gen. model de-biasing	$\mathcal{H}(p^\pi) \coloneqq -\mathop{\mathbb{E}}_{x\sim p^\pi}[\log p^\pi(x)]$	✗	✓	✓
Risk-averse optimization	$\mathrm{CVaR}^r_\beta(p^\pi) \coloneqq \mathop{\mathbb{E}}_{x\sim p^\pi}[r(x) \mid r(x) \le \mathrm{q}^r_\beta(p^\pi)]$	✗	✓	✓
	$\mathbb{E}_{x\sim p^\pi}[r(x)] - \mathbb{V}\mathrm{ar}(p^\pi)$	✗	✗	✓
Novelty-seeking optimization	$\mathrm{SQ}^r_\beta(p^\pi) \coloneqq \mathop{\mathbb{E}}_{x\sim p^\pi}[r(x) \mid r(x) \ge \mathrm{q}^r_\beta(p^\pi)]$	✗	✗	✓
Optimal Experiment Design	$\mathrm{s}\left(\mathop{\mathbb{E}}_{x\sim p^\pi}[\Phi(x)\Phi(x)^\top - \lambda\mathbb{I}]\right)$ $\mathrm{s}(\cdot) \in \{\log\det(\cdot), -\mathrm{Tr}(\cdot)^{-1}, -\lambda_{max}(\cdot)\}$	✗	✓	✓
Diverse modes discovery	$-\mathop{\mathbb{E}}_{z}[D_{KL}(p^{\pi,z} \Vert \mathop{\mathbb{E}}_{k} p^{\pi,k})]$	✗	✗	✓
Log-Barrier Constrained Generation	$\mathbb{E}_{x\sim p^\pi}[r(x)] - \beta\log(\langle p^\pi, c\rangle - C)$	✗	✓	✓
Kullback–Leibler divergence [14, 55]	$D_{KL}(p^\pi \Vert p^{pre}) = \int p^\pi(x)\log\frac{p^\pi(x)}{p^{pre}(x)}\,dx$	✓	✓	✓
Rényi divergences	$D_\beta(p^\pi \Vert p^{pre}) \coloneqq \frac{1}{\beta-1}\log\int (p^\pi(x))^\beta (p^{pre})^{1-\beta}\,dx$	✗	✗	✓
Optimal Transport distances	$W_p(p^\pi \Vert p^{pre}) \coloneqq \inf_{\gamma\in\Gamma(p^\pi,p^{pre})} \mathop{\mathbb{E}}_{(x,y)\sim\gamma}[d(x,y)^p]^{\frac{1}{p}}$	✗	✗	✓
Maximum Mean Discrepancy	$\mathrm{MMD}_k(p^\pi \Vert p^{pre}) \coloneqq \Vert\mu_{p^\pi} - \mu_{p^{pre}}\Vert,\ \mu_p \coloneqq \mathop{\mathbb{E}}_{x\sim p}[k(x,\cdot)]$	✗	✓	✓

Flow Density Control: Theoretical Guarantees

Distributional Generative Optimization

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

Flow Density Control: Theoretical Guarantees

Distributional Generative Optimization

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

$$\mathcal{G}(p_1^\pi)$$

Flow Density Control: Theoretical Guarantees

Distributional Generative Optimization

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

$$\mathcal{G}(p_1^\pi)$$

concave, idealized assumptions, K iterations:

$$\mathcal{G}(p_1^*) - \mathcal{G}(p_1^\pi) \leq \frac{L-l}{K} D_{KL}(p_1^* \| p_1^{pre})$$

L-relative smoothness, l-relative strong convexity

(Finite-time convergence via Nesterov relative smoothness and strong convexity framework)

general (e.g., non-concave) assumptions:

$$p_1^k \to \tilde{p_1}$$

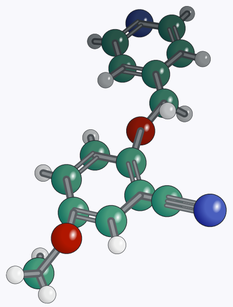
$\tilde{p_1}$ stationary point of $\mathcal{G}$.

(Asymptotic convergence via mirror flows)

Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning

[**Riccardo De Santi**, Marin Vlastelica, Ya-Ping Hsieh, Zebang Shen, Niao He, Andreas Krause]

Spotlight **NeurIPS 2025** **Oral** **at ICML 2025 BioGen Workshop**



Efficient Tail-Aware Generative Optimization via Flow Model Fine-Tuning

[Zifan Wang, **Riccardo De Santi**, Xiaoyu Mo, and Michael M. Zavlanos, Andreas Krause, Karl H. Johansson]

ICML 2026

Provable Maximum Entropy Manifold Exploration via Diffusion Models

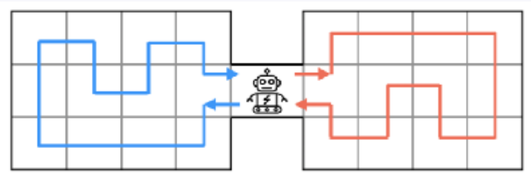
[**Riccardo De Santi***, Marin Vlastelica*, Ya-Ping Hsieh, Zebang Shen, Niao He, Andreas Krause]

ICML 2025

The Importance of Non-Markovianity in Maximum State Entropy Exploration

[Mirco Mutti*, **Riccardo De Santi***, and Marcello Restelli]

Outstanding Paper Award at **ICML 2022**



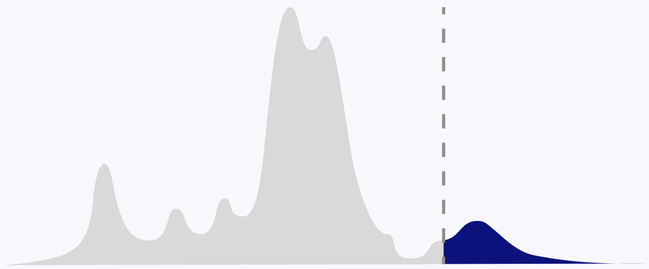
This talk:

Generative Discovery Beyond the Data

Part I

Rare and Diverse Generation

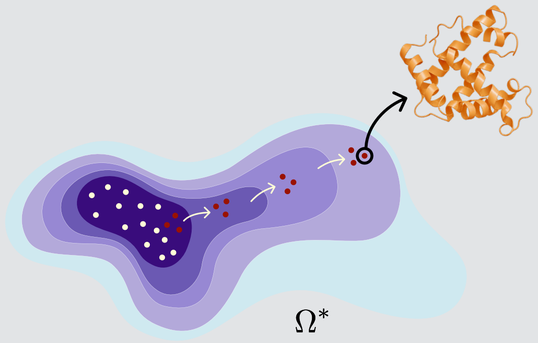
Which candidates matter?



Part II

Out-of-Distribution Flow Modeling

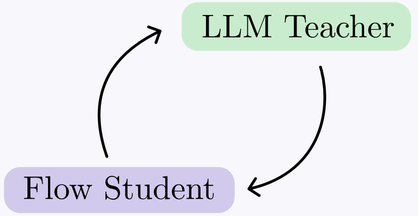
Which regions can we reach?



Part III

Sample-Efficient Discovery via LLM Feedback

What feedback guides us there?



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

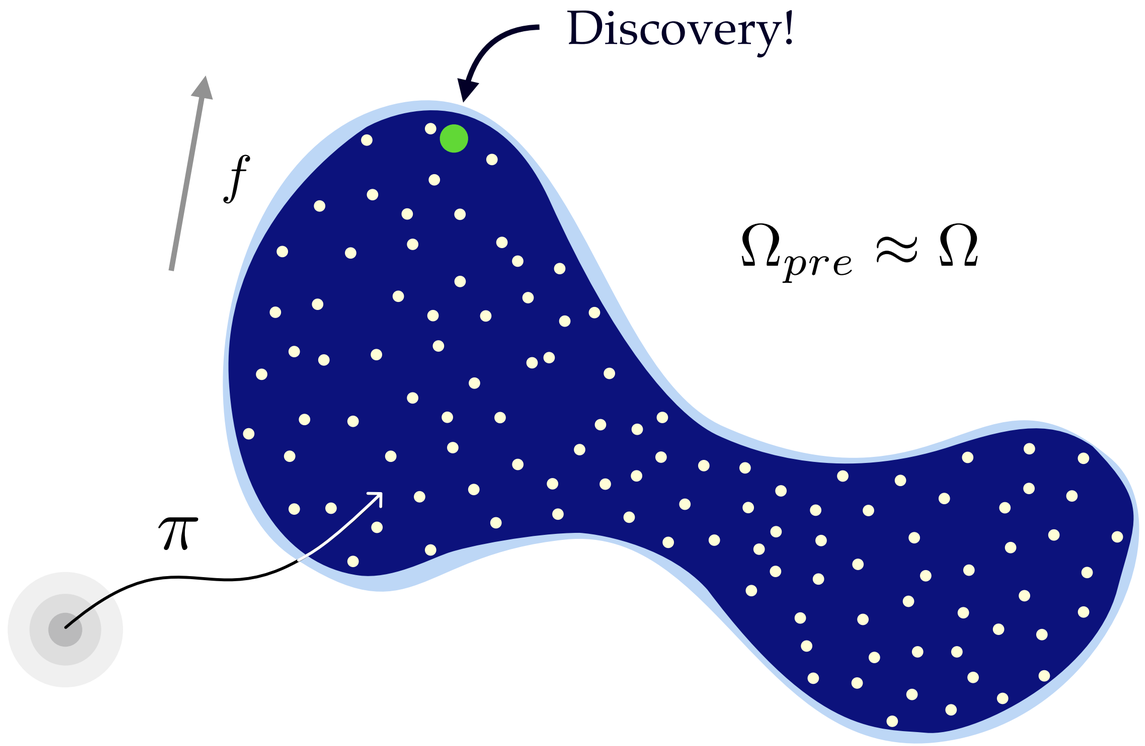
[**Riccardo De Santi**, Bruce Lee, Cristian Perez Jensen, Kimon Protopapas, Sophia Tang, Cheng-Hao Liu, Pranam Chatterjee, Yisong Yue, Andreas Krause]

NeurIPS 2026

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

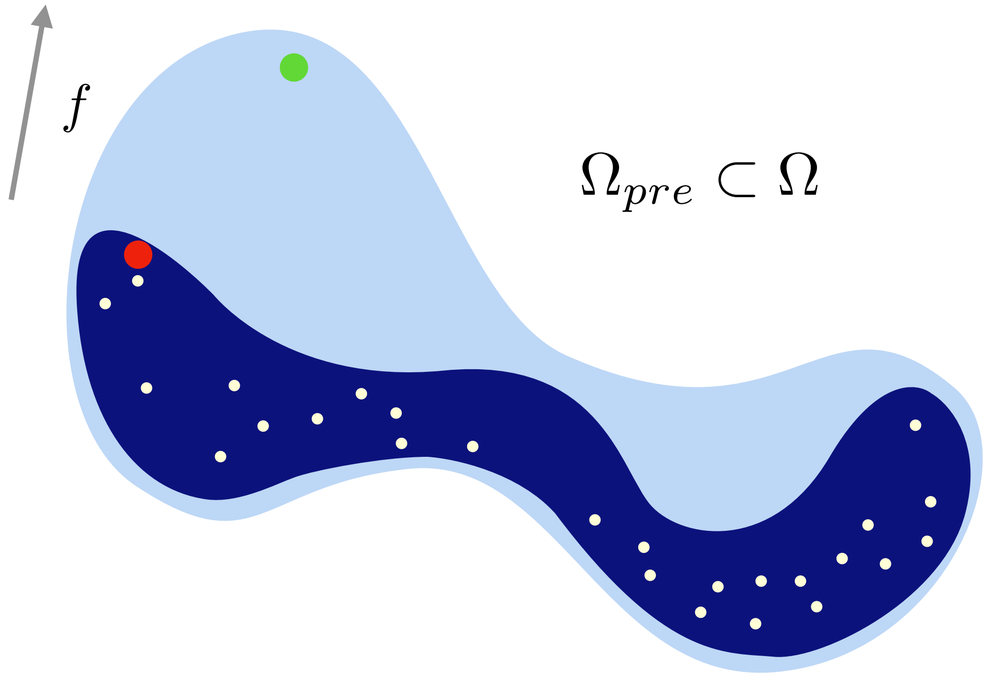
Discovery problem

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

The Reality: Generative Optimization Limitation

valid design space Ω

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



Generative models generalize locally over training data.

Especially in natural sciences:

E.g.,
- Chemistry
- Biology
- Physics

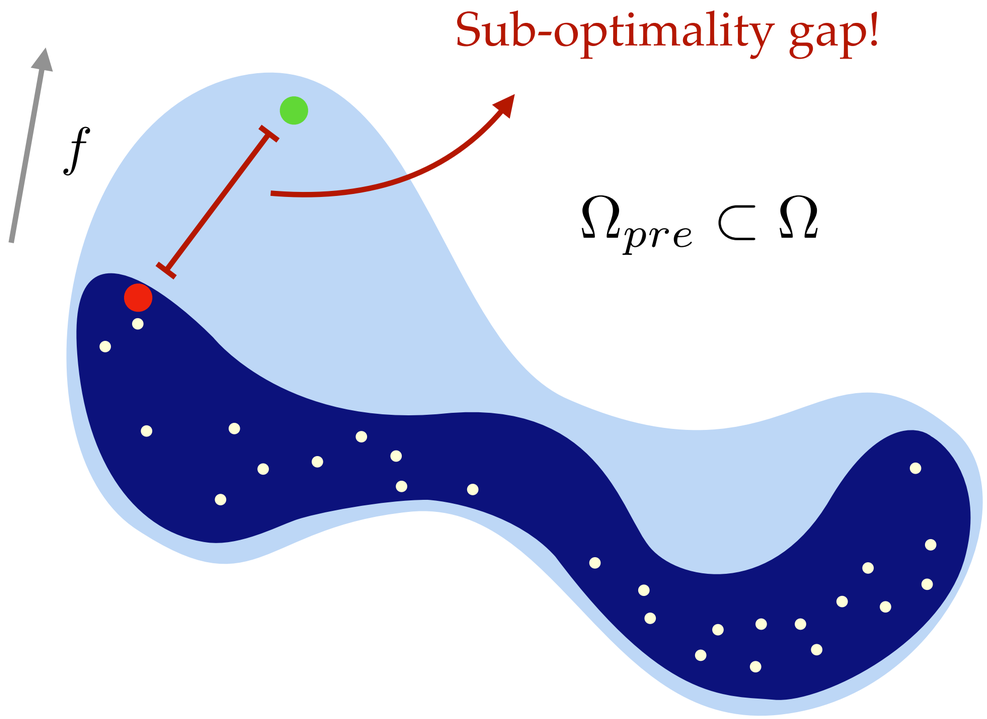
Discovery problem

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

The Reality: Generative Optimization Limitation

valid design space Ω

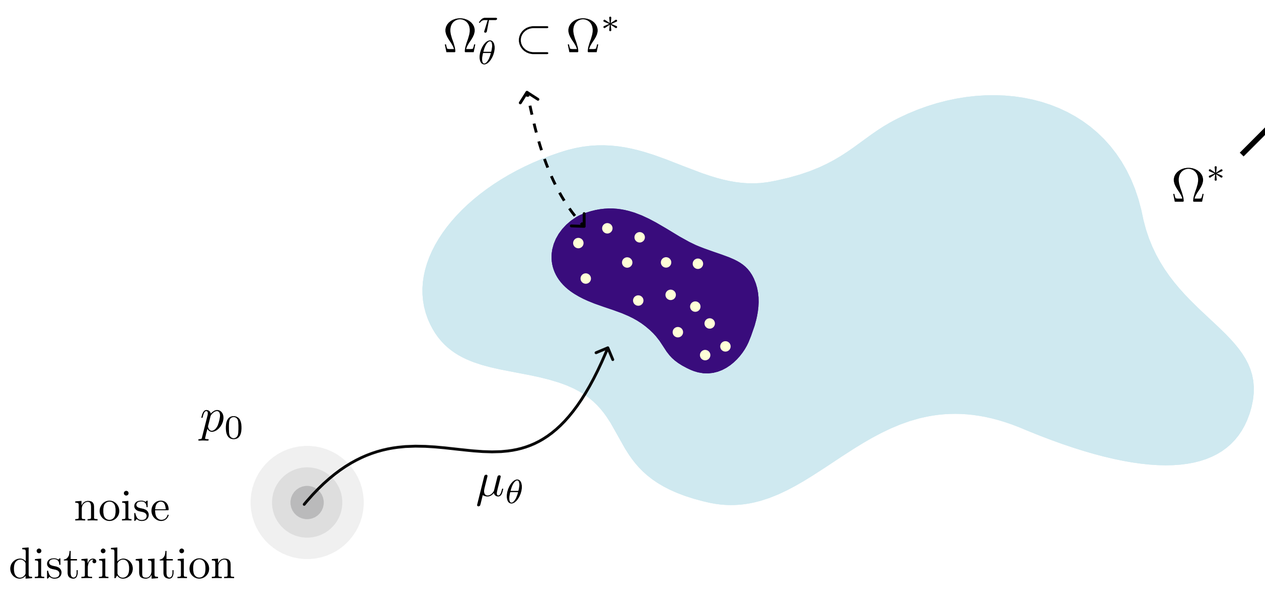
generable set of pre-trained gen. model Ω_{pre}



Discovery problem

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

Out-of-Distribution Flow Modeling



Valid design space

$\Omega^* := \{x \in \mathcal{X} : v(x) = 1\}$

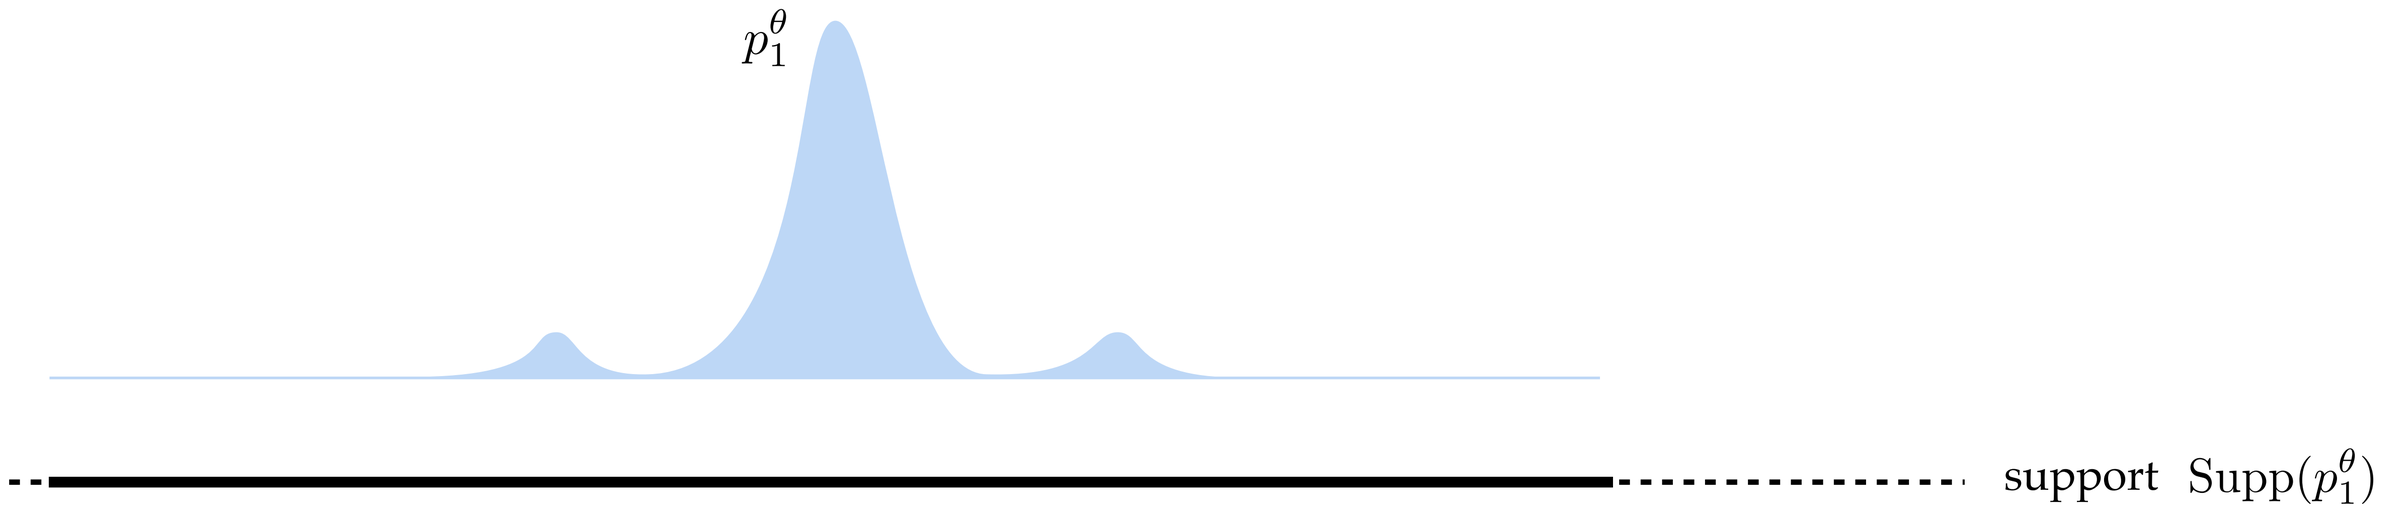
Ω^*

Unknown laws of nature
(e.g., chemistry, biology)

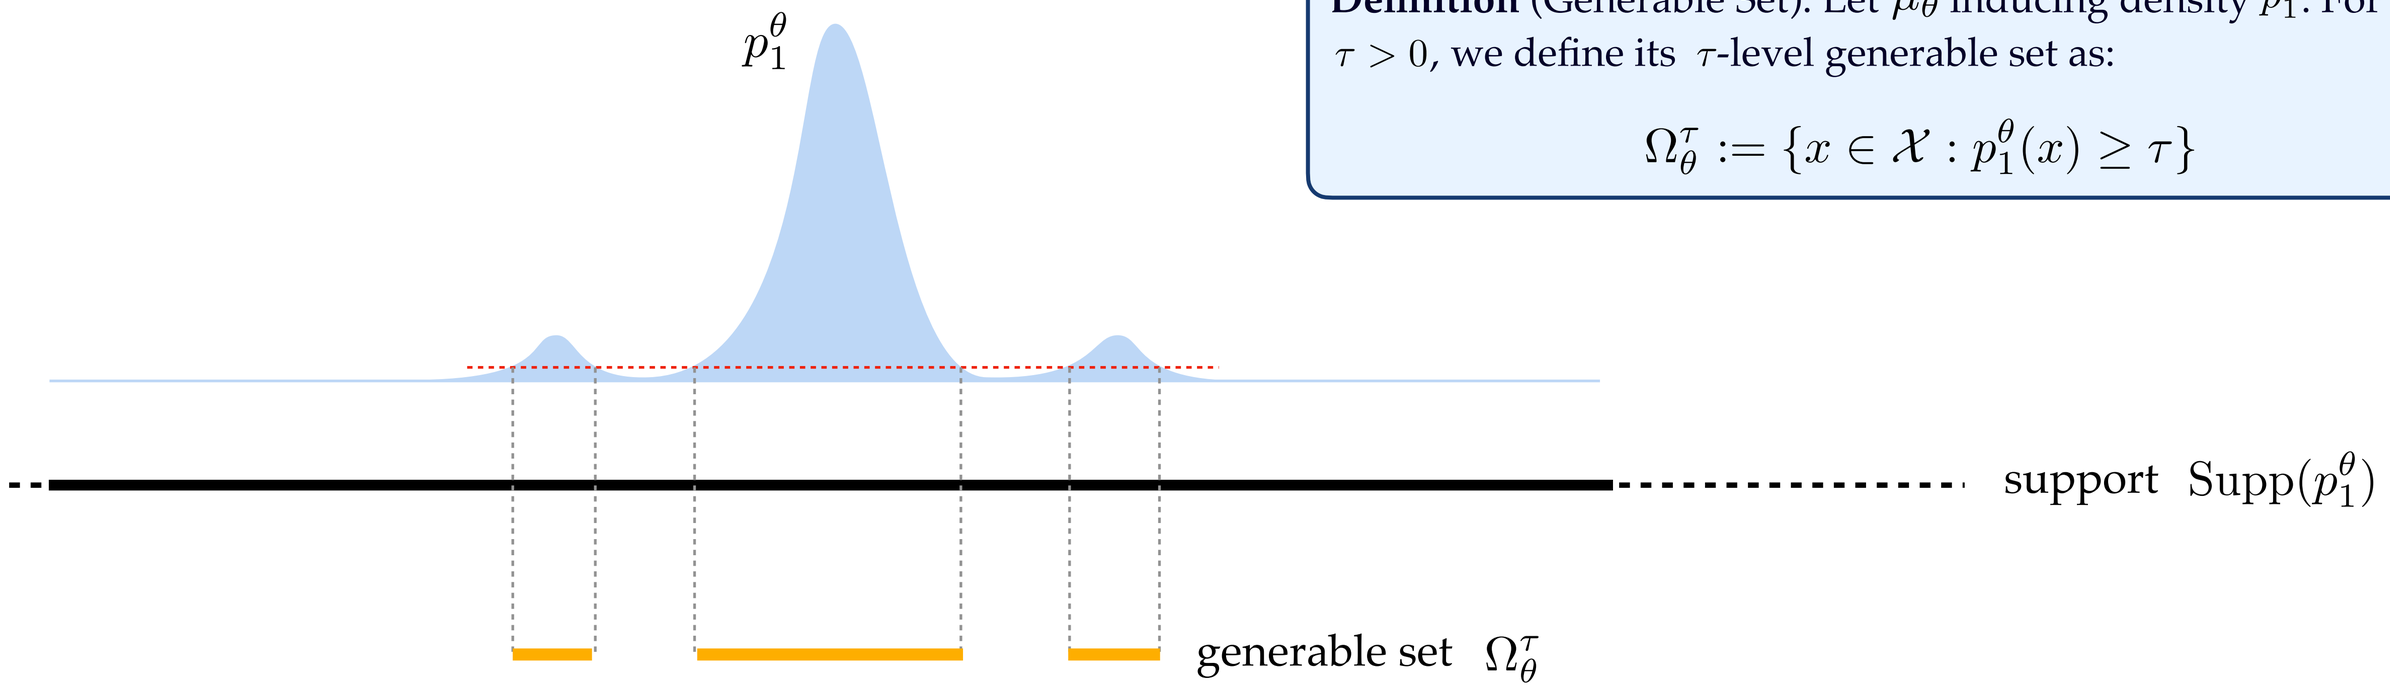
Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

Generable Set vs Support



Generable Set vs Support

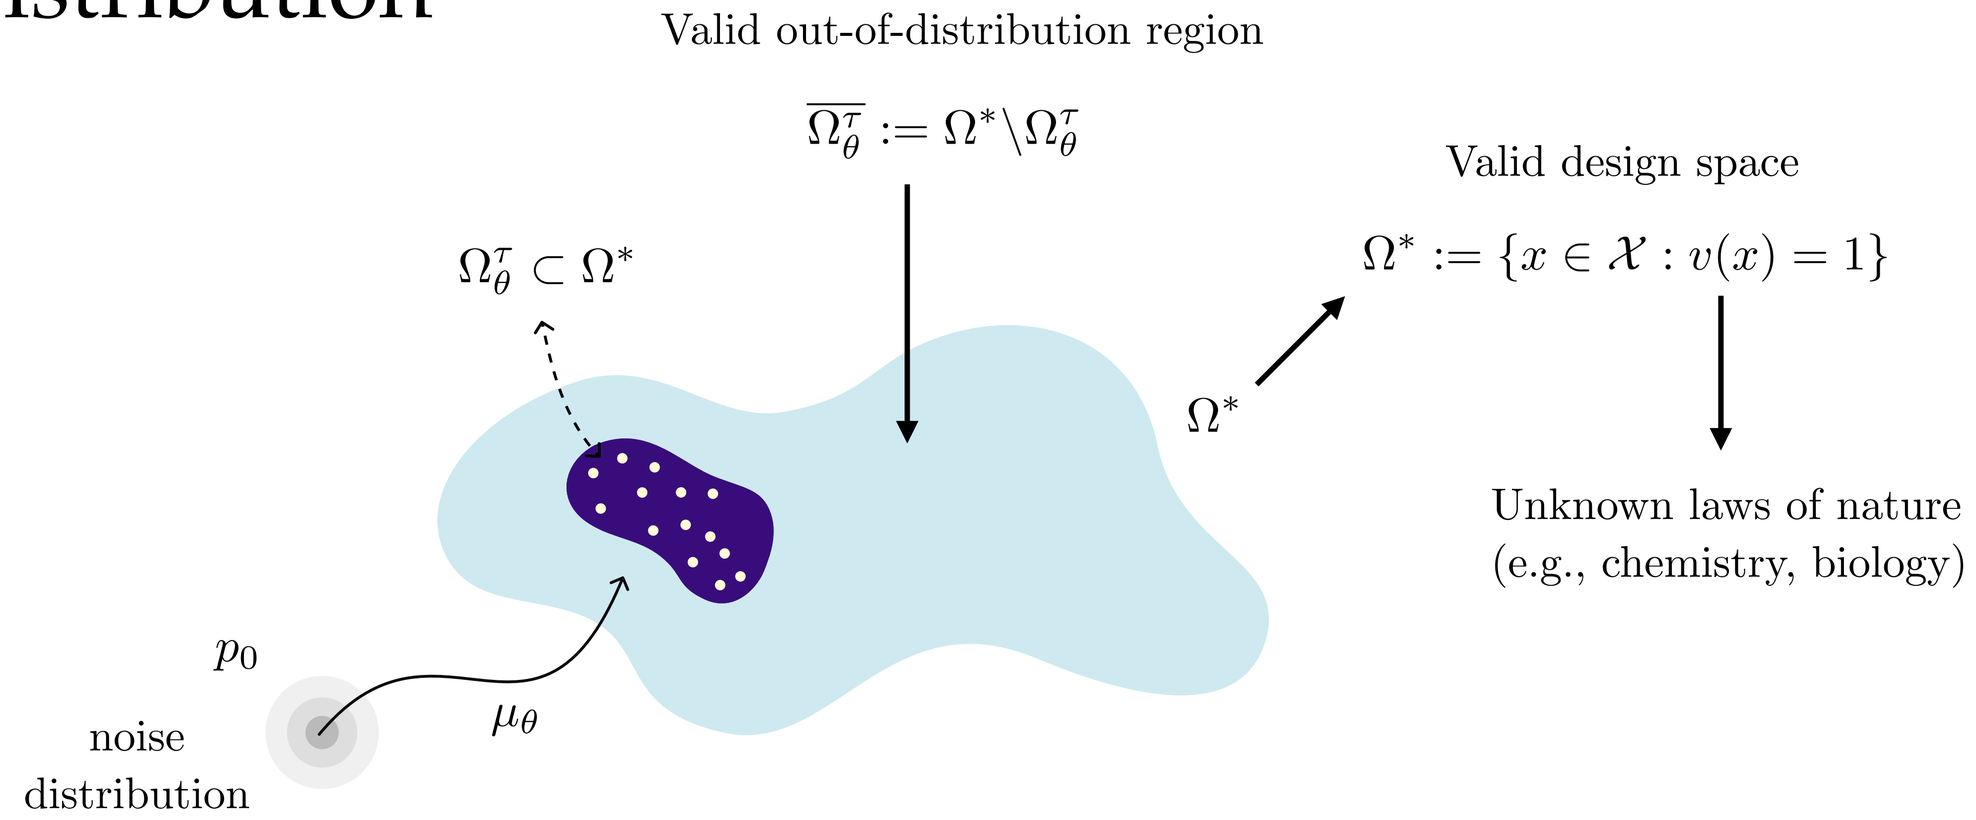


Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

generable set $\approx$ finite-sample coverage

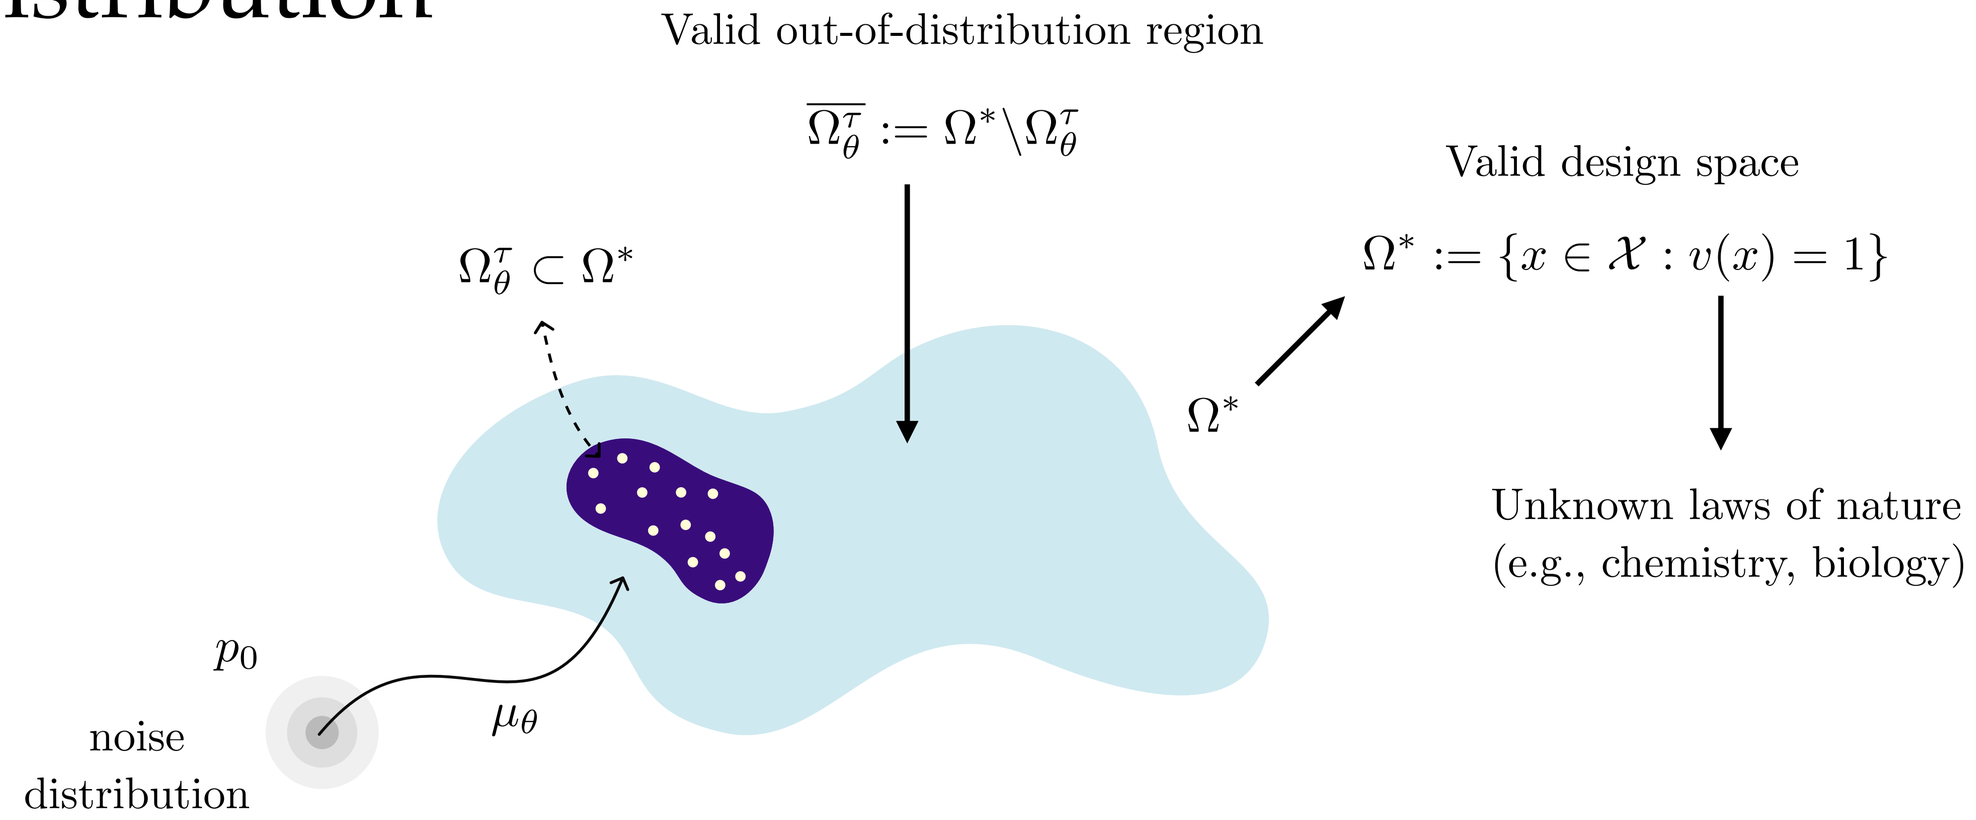
Beyond Distribution Matching



Generable Set Expansion:

learn θ s.t. $\Omega_\theta^\tau \approx \Omega^*$

Beyond Distribution Matching

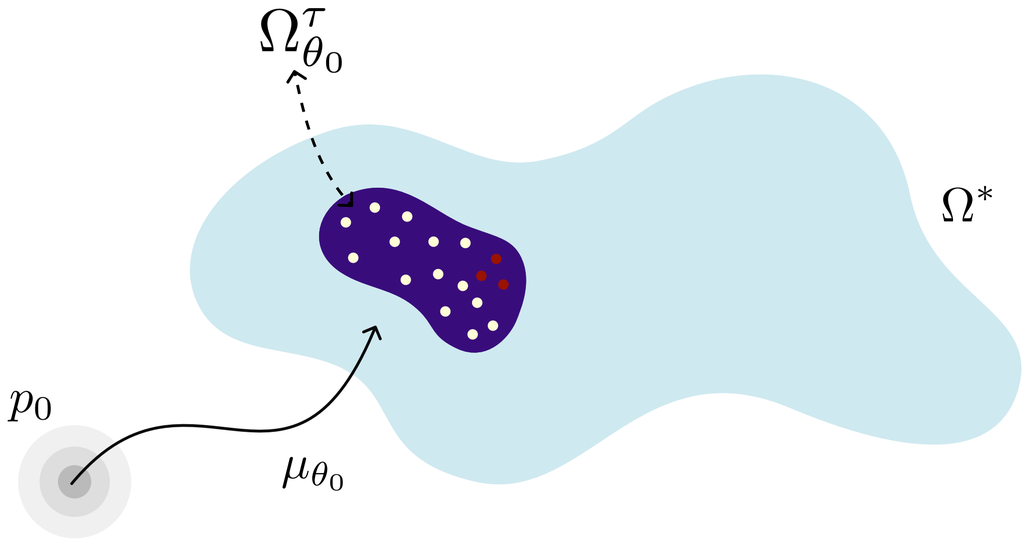


Requires acquiring information about Ω^* beyond the data

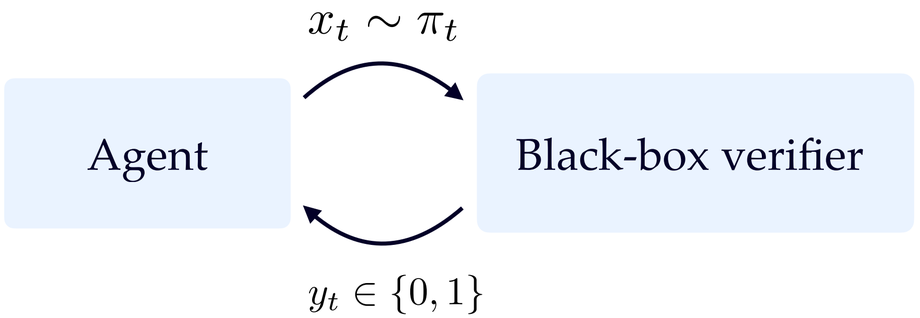
Generable Set Expansion:

learn θ s.t. $\Omega_\theta^\tau \approx \Omega^*$

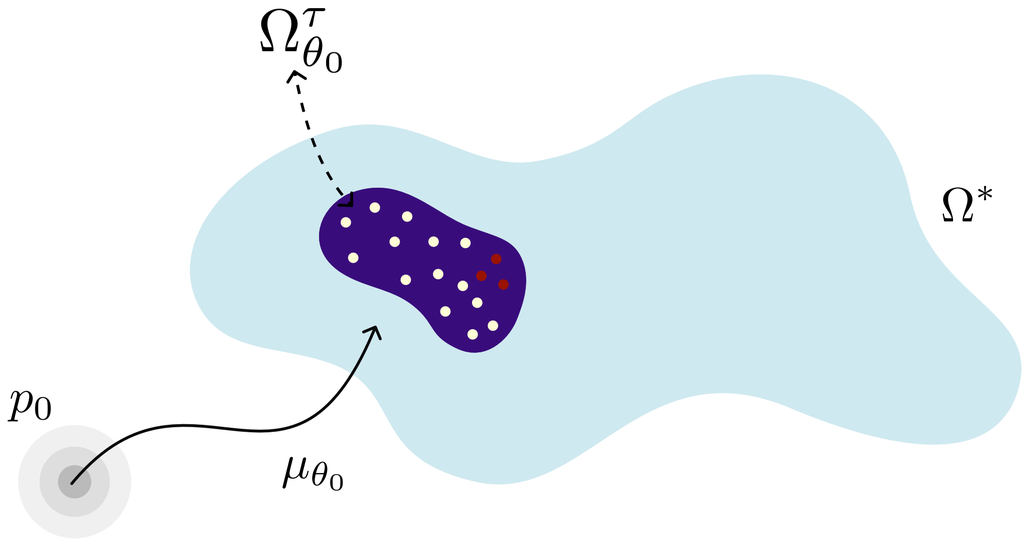
Generable Set Expansion for OOD Flow Modeling



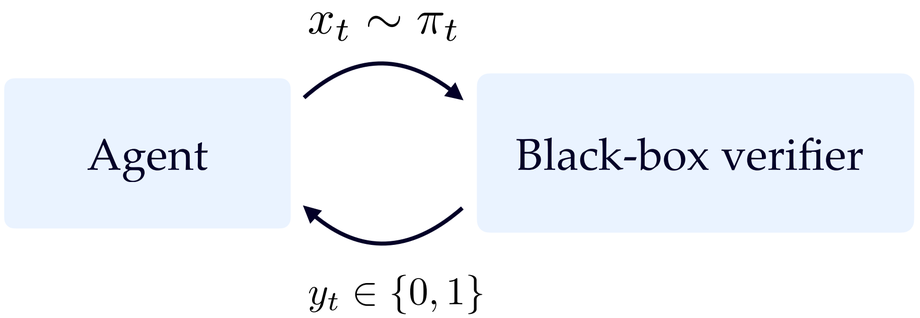
Black-box noisy binary verifier feedback



Generable Set Expansion for OOD Flow Modeling



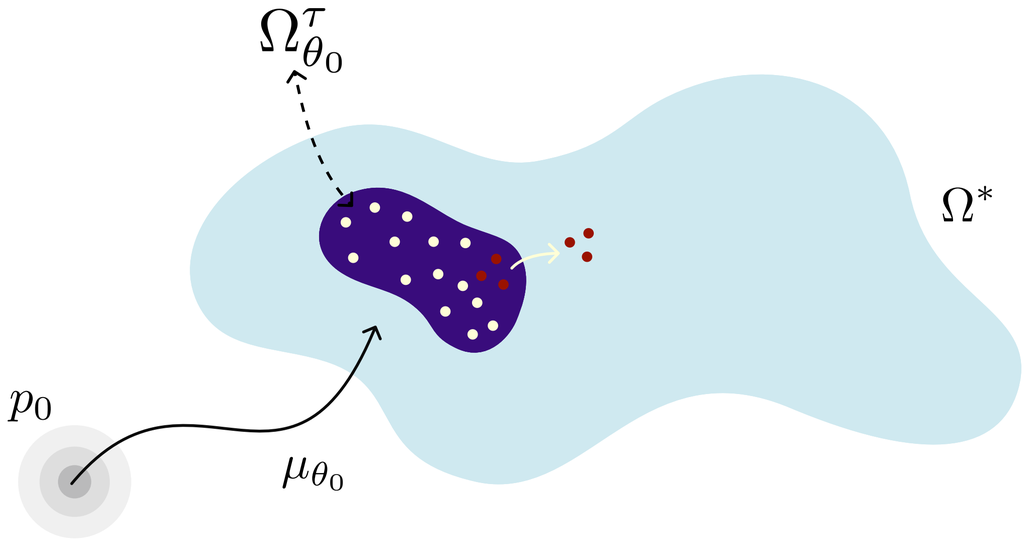
Black-box noisy binary verifier feedback



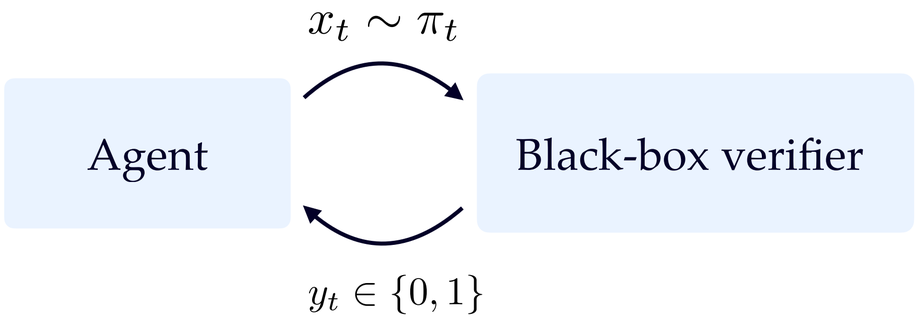
$\theta^{\text{pre}} = \theta_0$

$\Omega^{\tau}_{\theta_0}$

Generable Set Expansion for OOD Flow Modeling



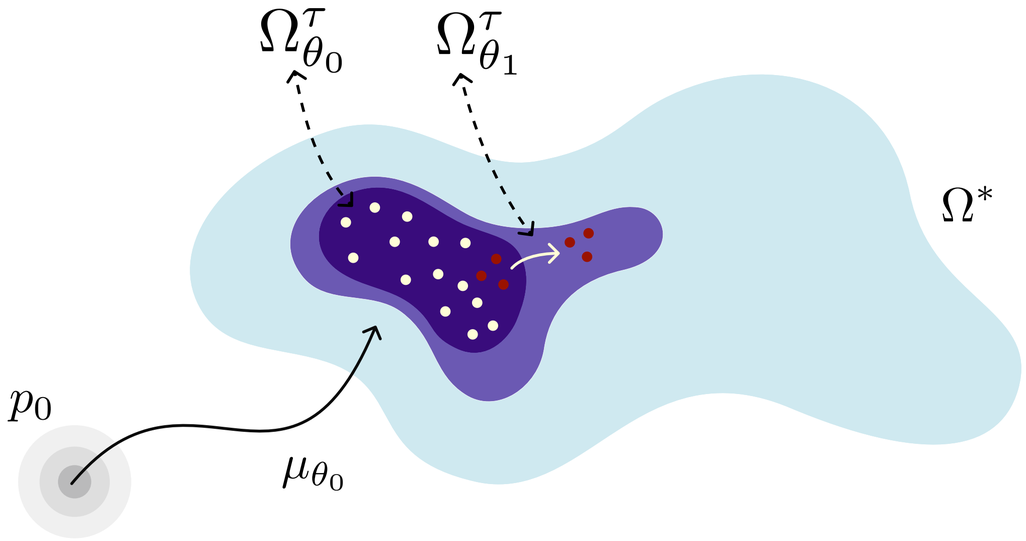
Black-box noisy binary verifier feedback



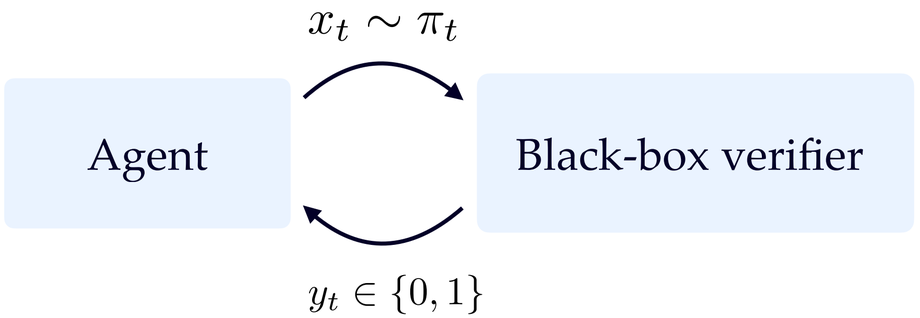
$\theta^{\text{pre}} = \theta_0$ $\Omega^{\tau}_{\theta_0}$

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

Generable Set Expansion for OOD Flow Modeling



Black-box noisy binary verifier feedback

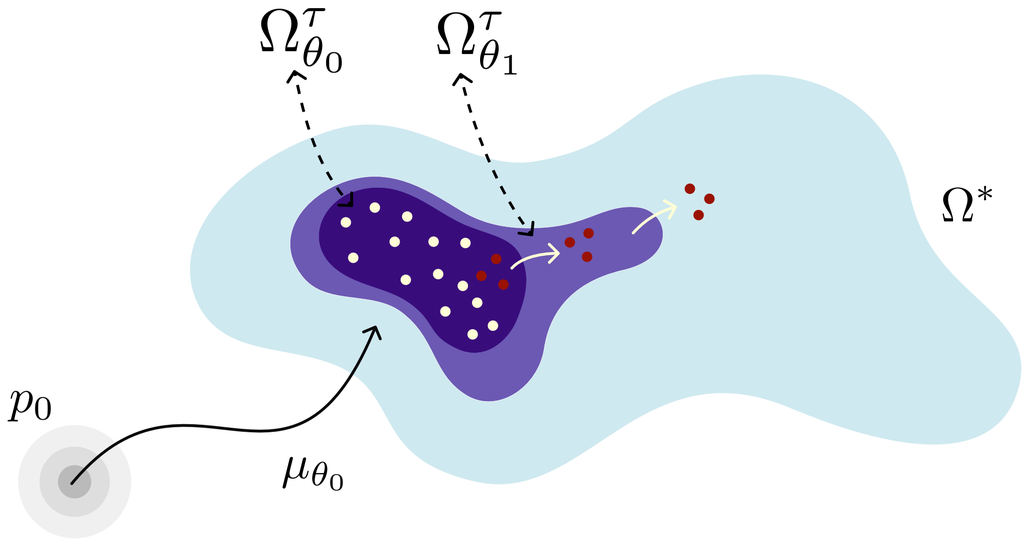


$$\theta^{\text{pre}} = \theta_0 \to \theta_1$$

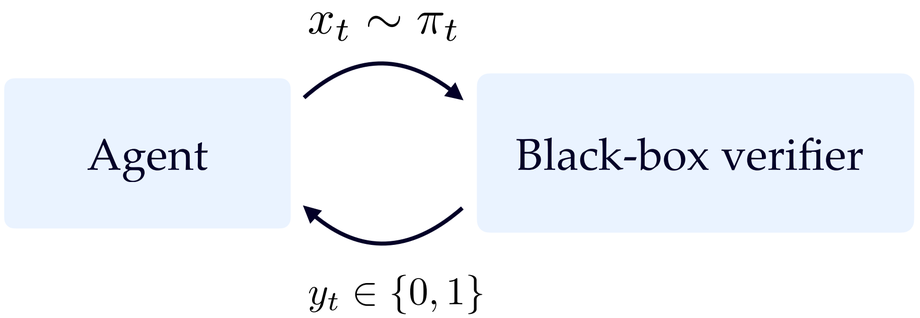
$$\Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1}$$

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

Generable Set Expansion for OOD Flow Modeling



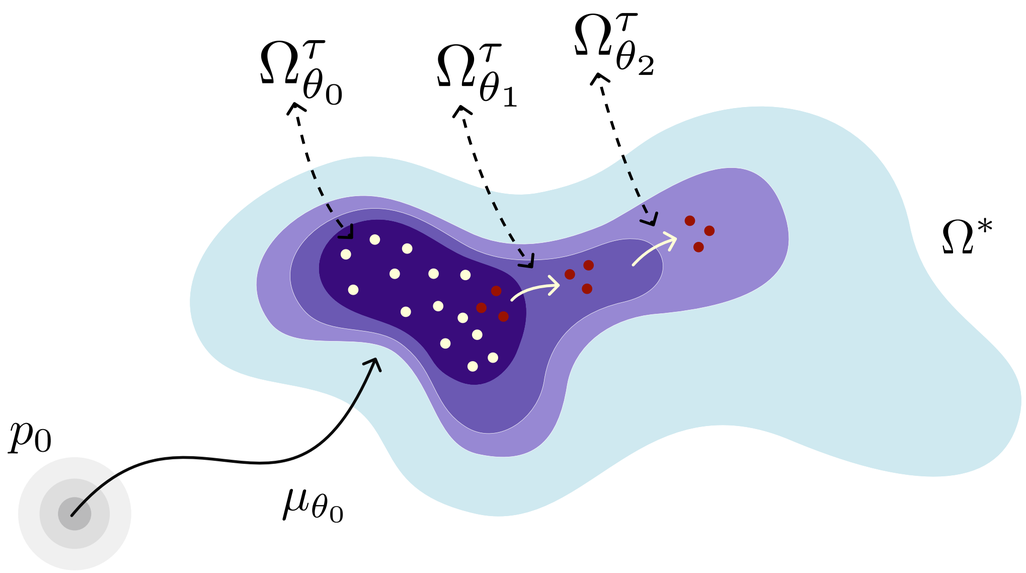
Black-box noisy binary verifier feedback



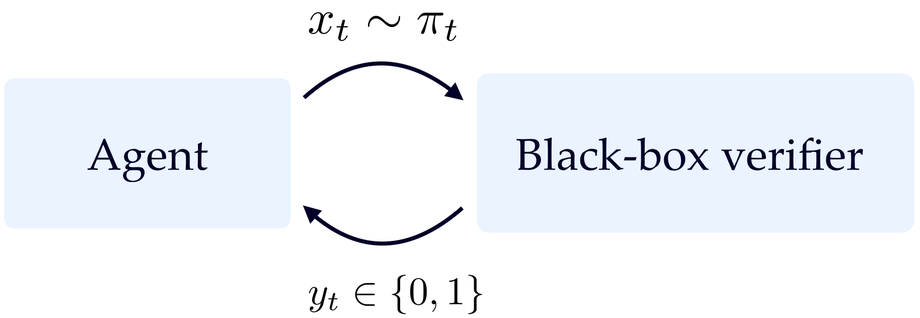
$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq$$

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

Generable Set Expansion for OOD Flow Modeling



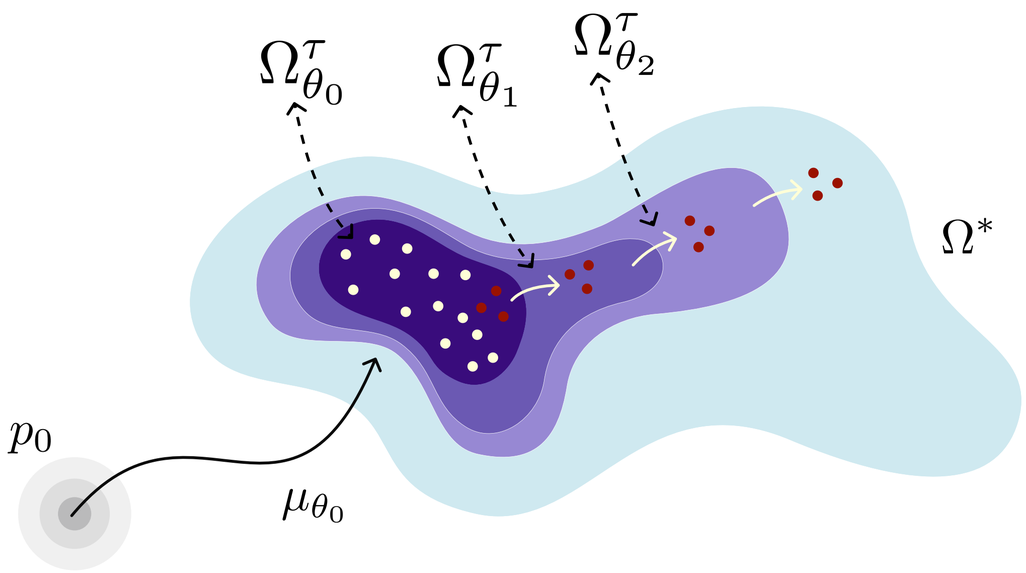
Black-box noisy binary verifier feedback



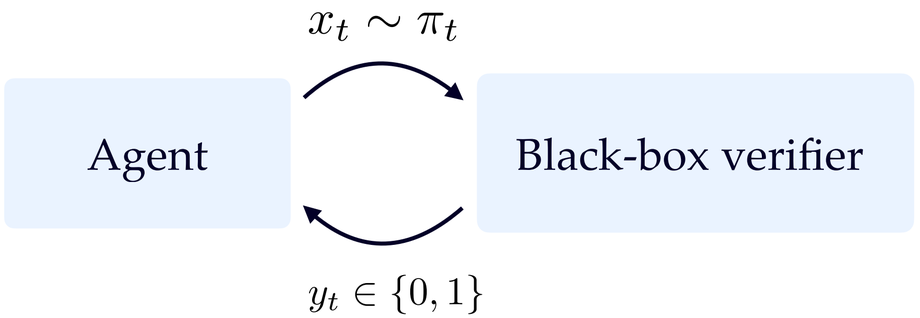
$$\theta^{\mathrm{pre}} = \theta_0 \to \theta_1 \to \cdots \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots$$

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

Generable Set Expansion for OOD Flow Modeling



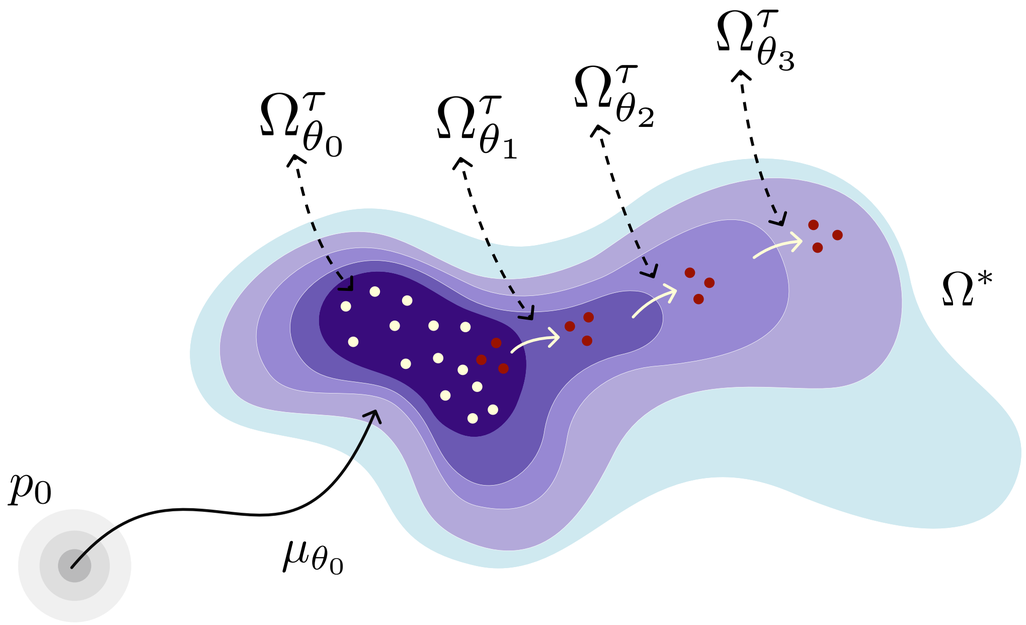
Black-box noisy binary verifier feedback



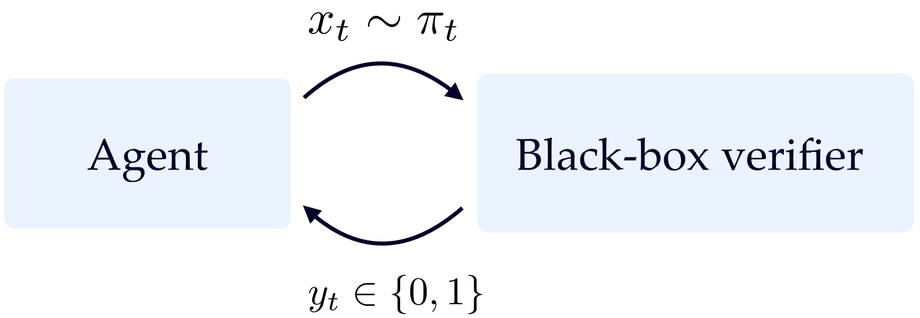
$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq$$

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

Generable Set Expansion for OOD Flow Modeling



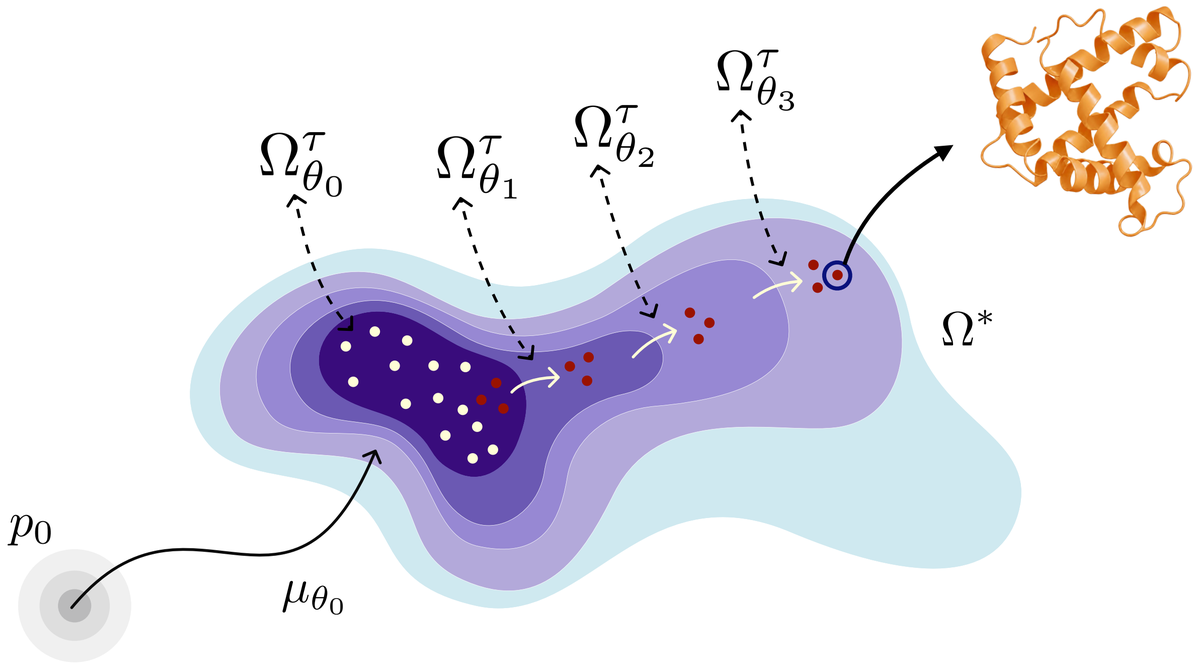
Black-box noisy binary verifier feedback



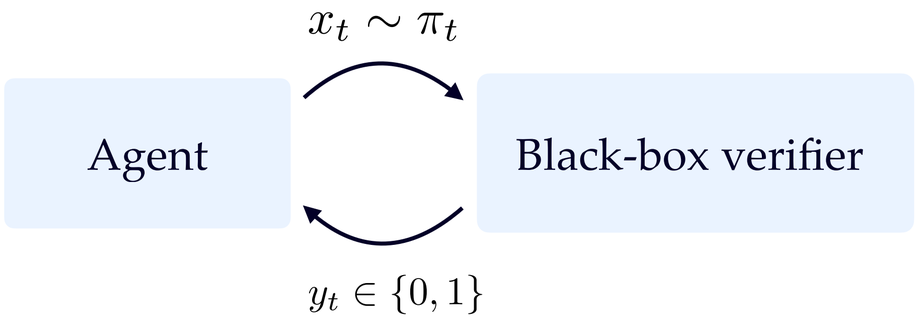
$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

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

Generable Set Expansion for OOD Flow Modeling



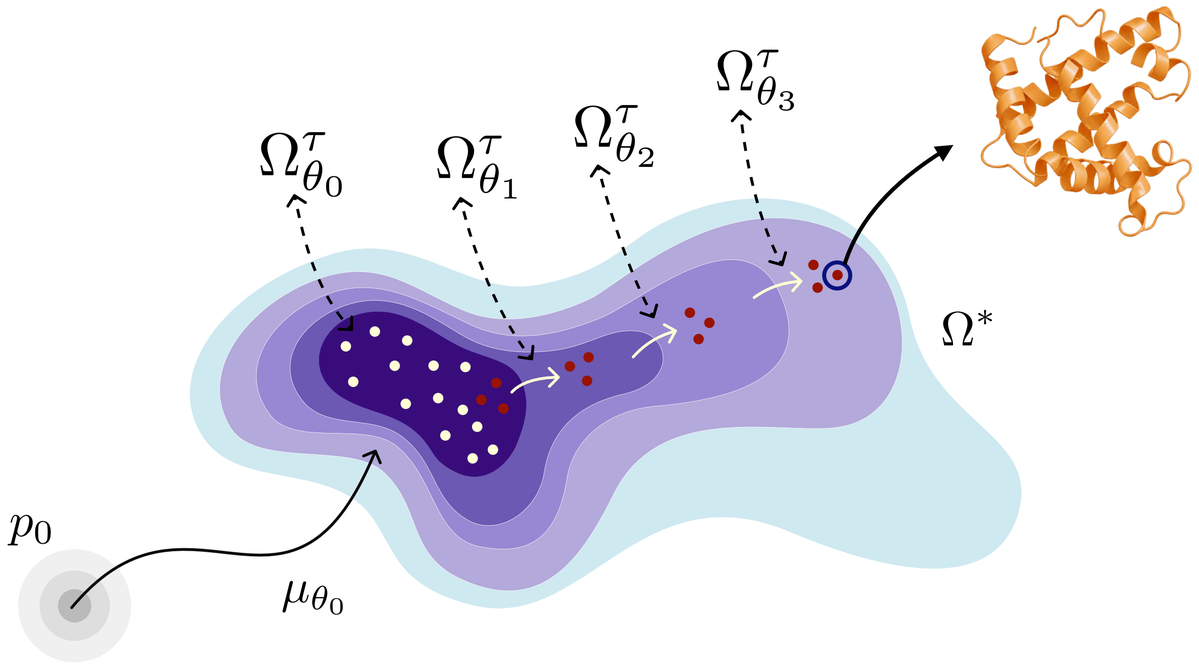
Black-box noisy binary verifier feedback



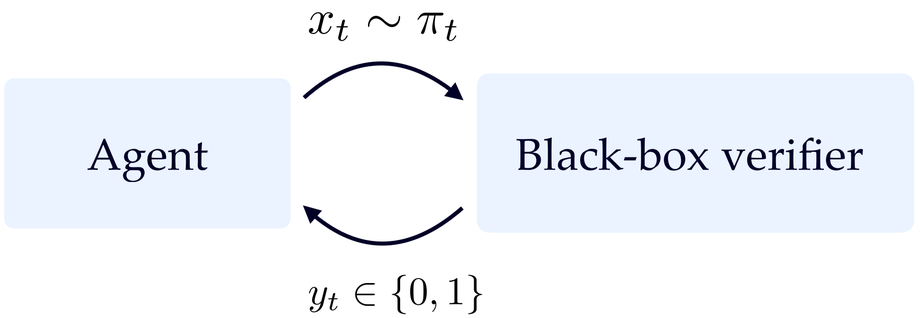
$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

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

Generable Set Expansion for OOD Flow Modeling



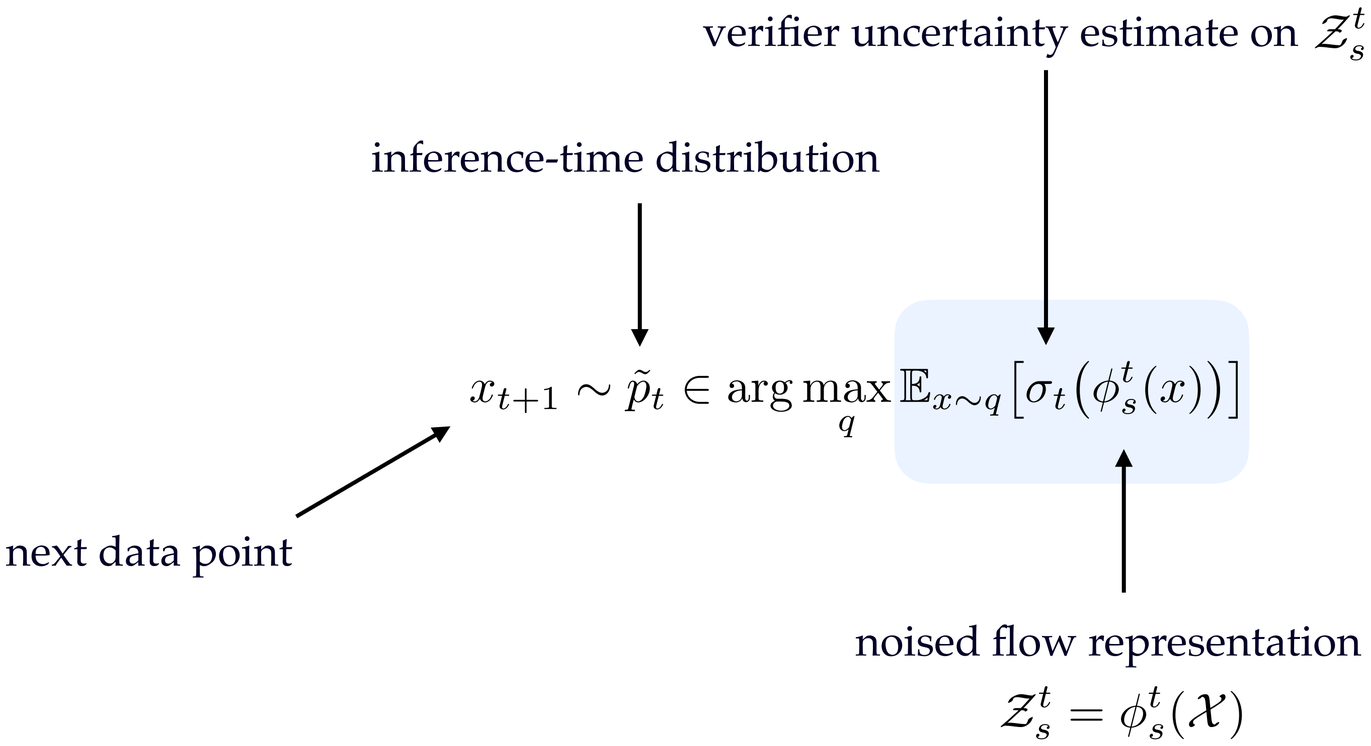
Black-box noisy binary verifier feedback



$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T, \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T} \overset{?}{\approx} \Omega^{\star}$$

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

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



generates data with **high verifier (i.e., laws of nature) uncertainty**

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

verifier uncertainty estimate on $\mathcal{Z}_s^t$

inference-time distribution

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

next data point

noised flow representation
$\mathcal{Z}_s^t = \phi_s^t(\mathcal{X})$

current-model (i.e., prior) regularization

generates data with **high verifier (i.e., laws of nature) uncertainty**, yet with **sufficient validity prior**

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

verifier uncertainty estimate on $\mathcal{Z}_s^t$

inference-time distribution

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

next data point

noised flow representation
$\mathcal{Z}_s^t = \phi_s^t(\mathcal{X})$

current-model (i.e., prior) regularization

generates data with **high verifier (i.e., laws of nature) uncertainty**, yet with **sufficient validity prior**

⟶ highly informative **frontier samples**

Theoretical Guarantees for Out-of-Distribution Flow Modeling

(Offline)
Distribution matching

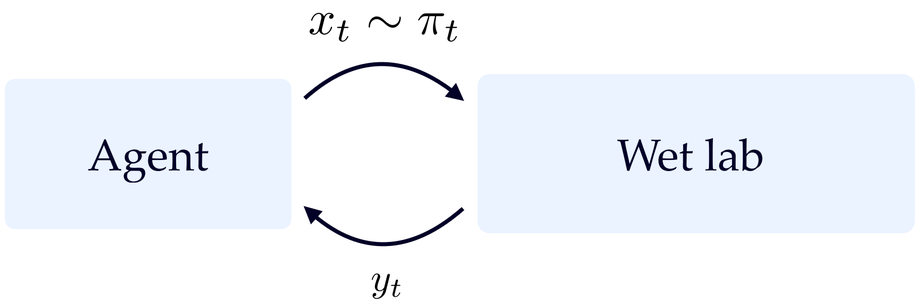
$p_1^\theta \approx p_{data}$

(Our work)

(Interactive)
Design space coverage

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

Discovery requires **coverage guarantees**, not only distributional fit.



Theoretical Guarantees for Out-of-Distribution Flow Modeling

Theorem (Informal, generable representation set covers reachable set). Fix $\epsilon > 0$ and an integer $H \geq 1$. Then, with

$$T^\star \gtrsim \left(\frac{\alpha \gamma^{\Omega_\star}_{HT^*}}{\epsilon} \right)^2$$

γ maximum information gain

More **complex validity laws** (of nature) **require more verifier queries**.

By running ActFlow, it holds with probability at least $1 - \delta$ that after T^* verified samples:

$$R^H_\epsilon(S_0) \subseteq \Omega^\tau_{T^\star}$$

generable representation set

H-step reachable set given recursive application of **one-step reachability operator** over flow representation:

Theoretical Guarantees for Out-of-Distribution Flow Modeling

Theorem (Informal, generable representation set covers reachable set). Fix $\epsilon > 0$ and an integer $H \geq 1$. Then, with

$$T^\star \gtrsim \left(\frac{\alpha \gamma^{\Omega_\star}_{HT^*}}{\epsilon} \right)^2$$

By running ActFlow, it holds with probability at least $1 - \delta$ that after T^* verified samples:

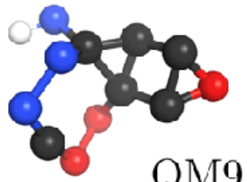
$$R^H_\epsilon(S_0) \subseteq \Omega^\tau_{T^\star}$$

Takeaway

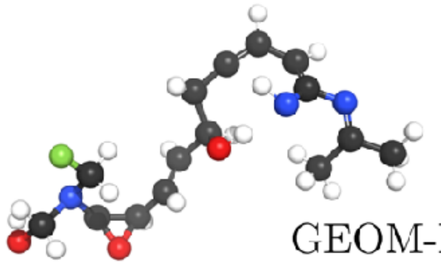
ActFlow provides first-of-their-kind finite-sample **coverage guarantees for out-of-distribution generative modeling**, under reachability assumptions.

Reachability + Sufficient Verifier Queries ⟶ **Design space coverage** $\Omega^\tau_\theta \approx \Omega^*$

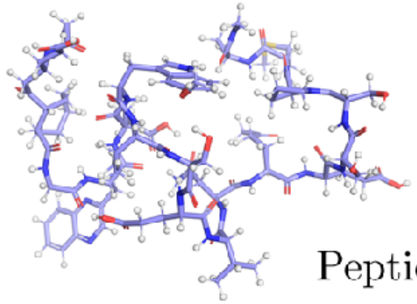
ActFlow with Chemical and Biological Data



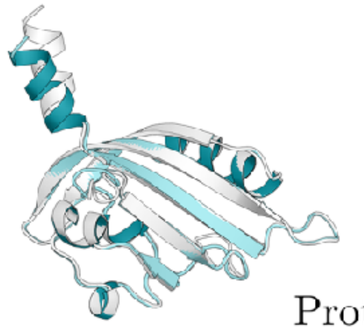
QM9 small
organic molecule



GEOM-Drugs
drug-like molecule

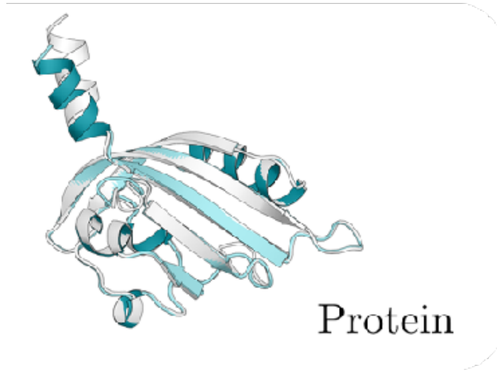


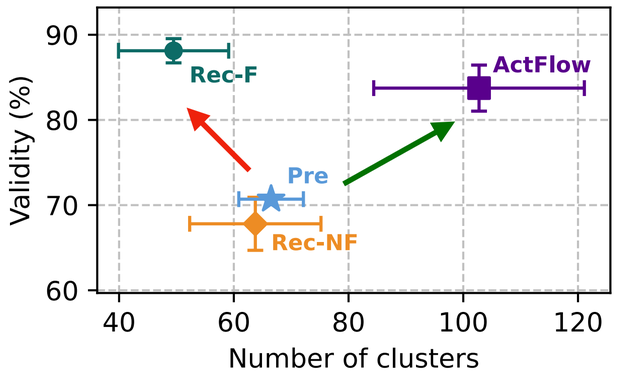
Peptide



Protein

ActFlow for Protein Design





Takeaway

ActFlow allows to **actively expand** pre-trained biochemical flow and diffusion models **over OOD regions**.

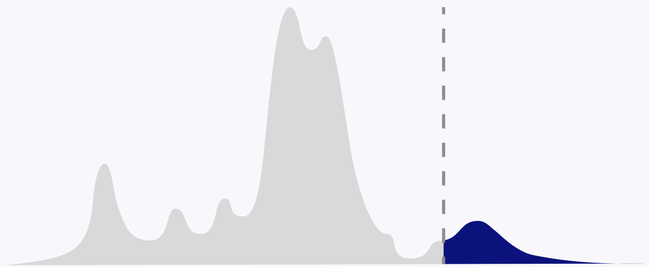
This talk:

Generative Discovery Beyond the Data

Part I

Rare and Diverse Generation

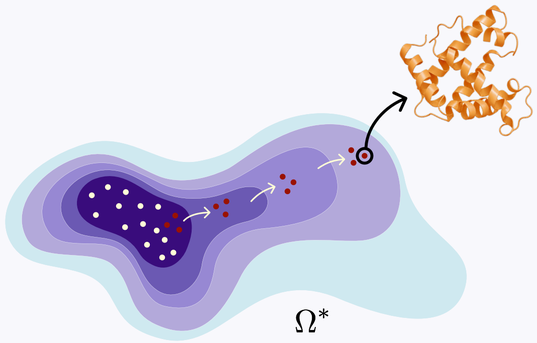
Which candidates matter?



Part II

Out-of-Distribution Flow Modeling

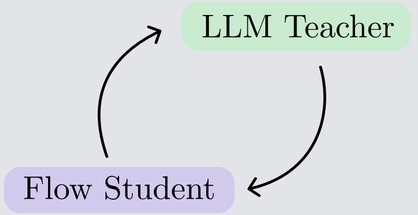
Which regions can we reach?



Part III

Sample-Efficient Discovery via LLM Feedback

What feedback guides us there?



Expansion for Discovery?

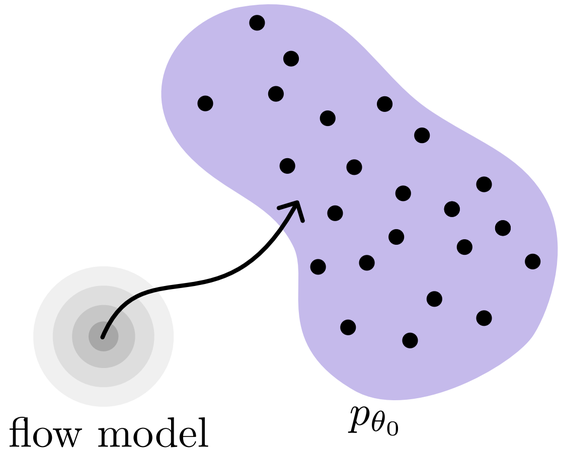
High-level discovery goal

'generate peptides with promising anticancer activity'

Expansion for Discovery?

High-level discovery goal

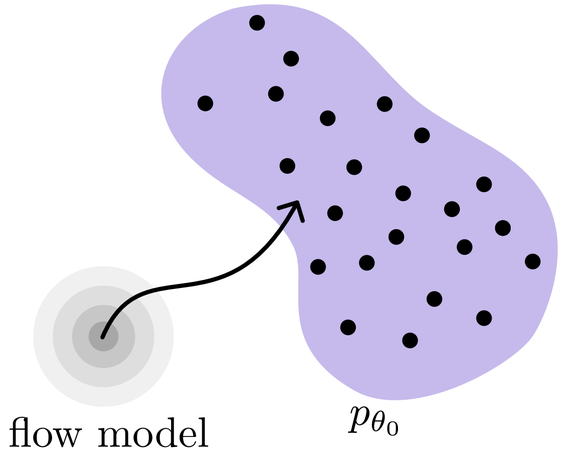
'generate peptides with promising anticancer activity'

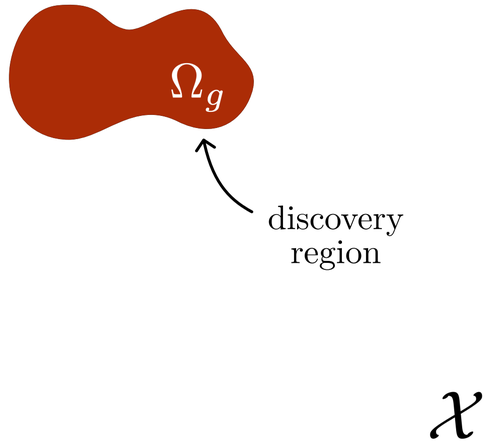


Expansion for Discovery?

High-level discovery goal

'generate peptides with promising anticancer activity'

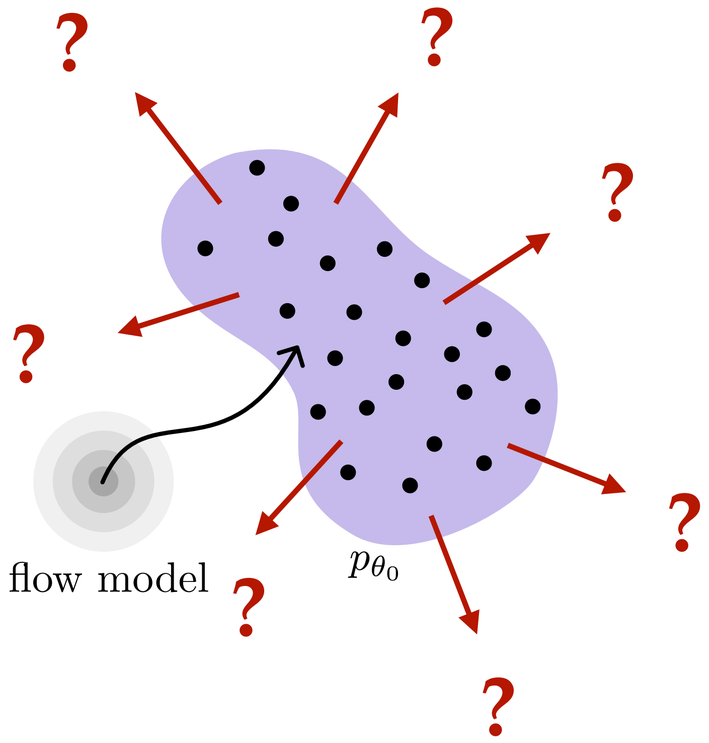


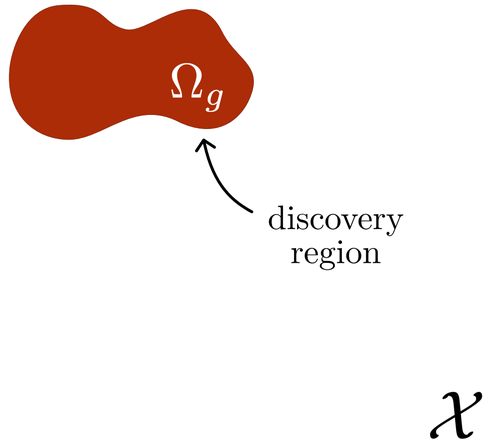


Goal-Directed Expansion for Discovery?

High-level discovery goal

'generate peptides with promising anticancer activity'



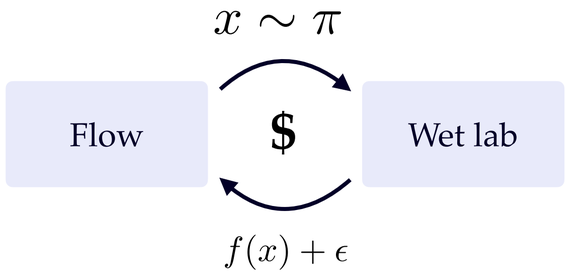


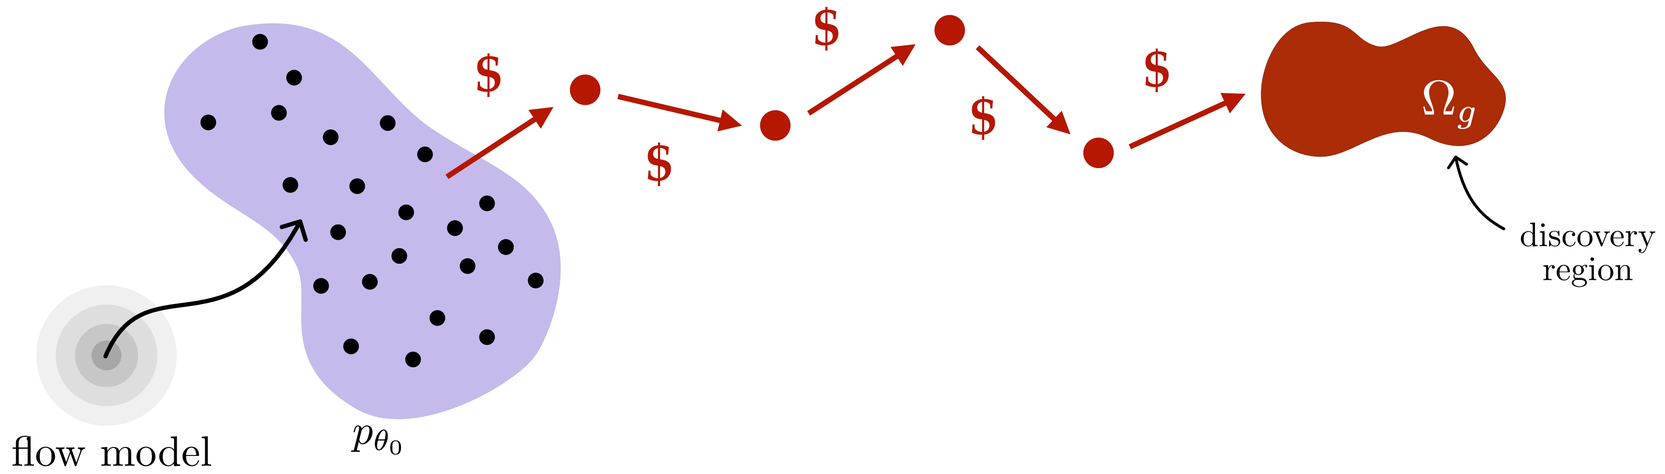
Very hard problem in high-dimensions!

Sample-Efficient and Goal-Directed Expansion for Discovery?

High-level discovery goal

'generate peptides with promising anticancer activity'

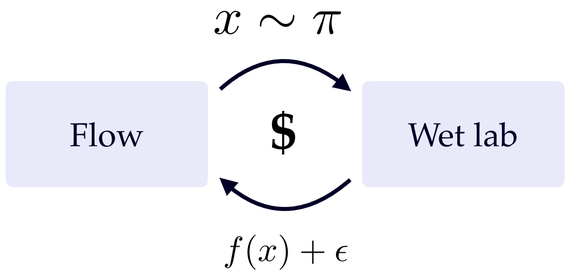


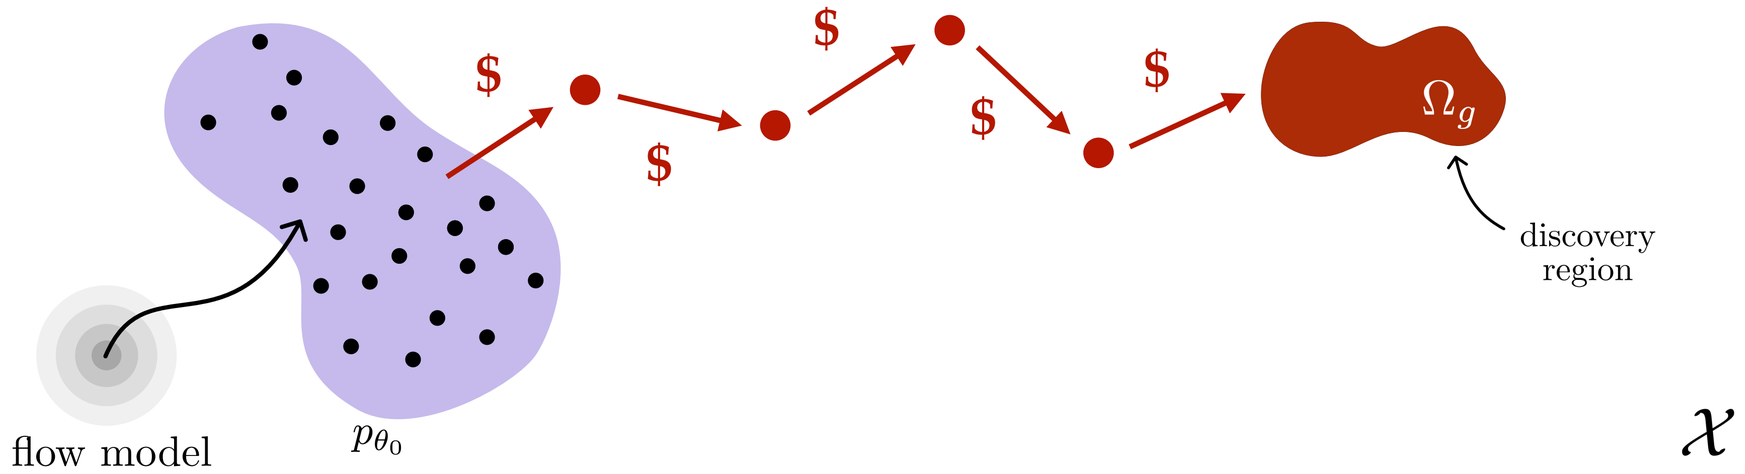


Sample-Efficient and Goal-Directed Expansion for Discovery?

High-level discovery goal

'generate peptides with promising anticancer activity'





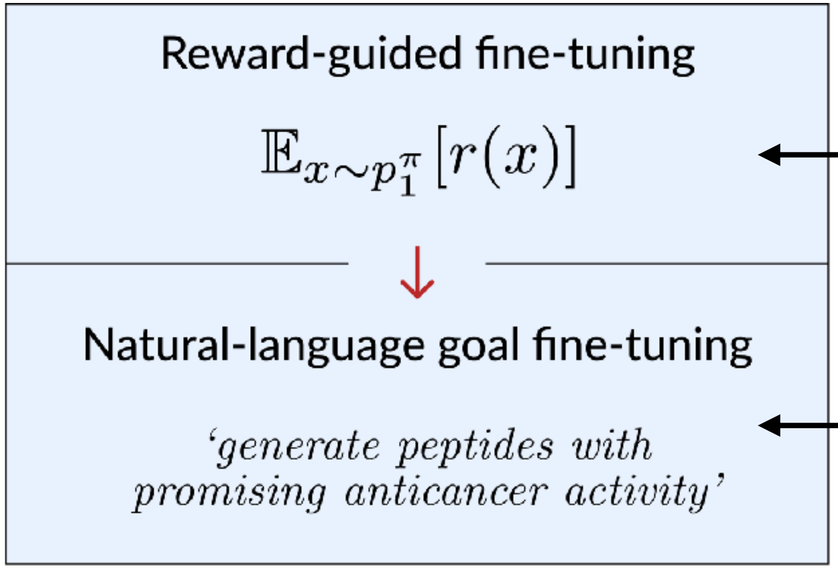
Can we exploit existing scientific knowledge in frontier LLMs for efficient discovery?

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

[**Riccardo De Santi***, Federico Di Gennaro*, Angelo Gnazzo*, Sophia Tang, Pranam Chatterjee, Fanny Yang, İsmail İlkan Ceylan, Michael M. Bronstein, Andreas Krause]

ArXiv Preprint

From Rewards to Natural-Language Goals



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

LFD (ours) (first method)

LFD is (to our knowledge) the **first method** to **adapt any scientific generator** to **natural-language goals**

Natural-language discovery goal

$g = \textit{'generate peptides with promising anticancer activity'}$ $\longrightarrow$

Discovery region

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

Natural-language discovery goal

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

Discovery region

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

Natural-language goal adaptation:

$$\text{update } \theta_0 \to \theta \;\; \text{such that} \;\; p_\theta(\Omega_g) \gg p_{\theta_0}(\Omega_g)$$

Natural-language discovery goal

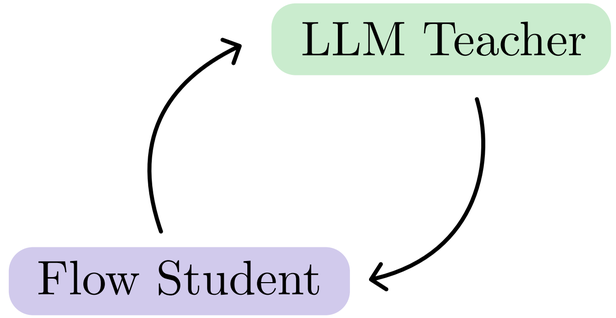
$g =$ *'generate peptides with promising anticancer activity'* $\longrightarrow$

Discovery region

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

Natural-language goal adaptation:

update $\theta_0 \to \theta$ such that $p_\theta(\Omega_g) \gg p_{\theta_0}(\Omega_g)$



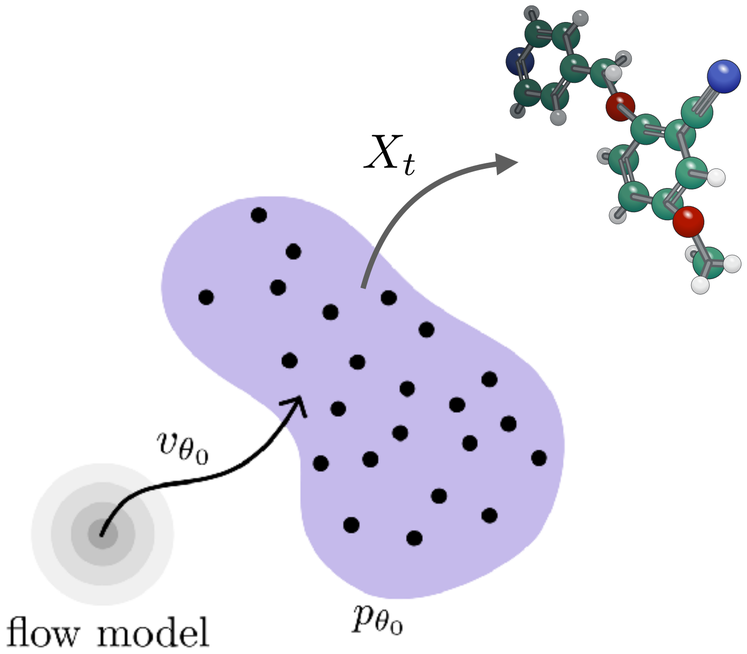
LLM-to-Flow Distillation: distill goal-related LLM scientific knowledge into a domain-specific generator.

Natural-language discovery goal

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

Discovery region

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



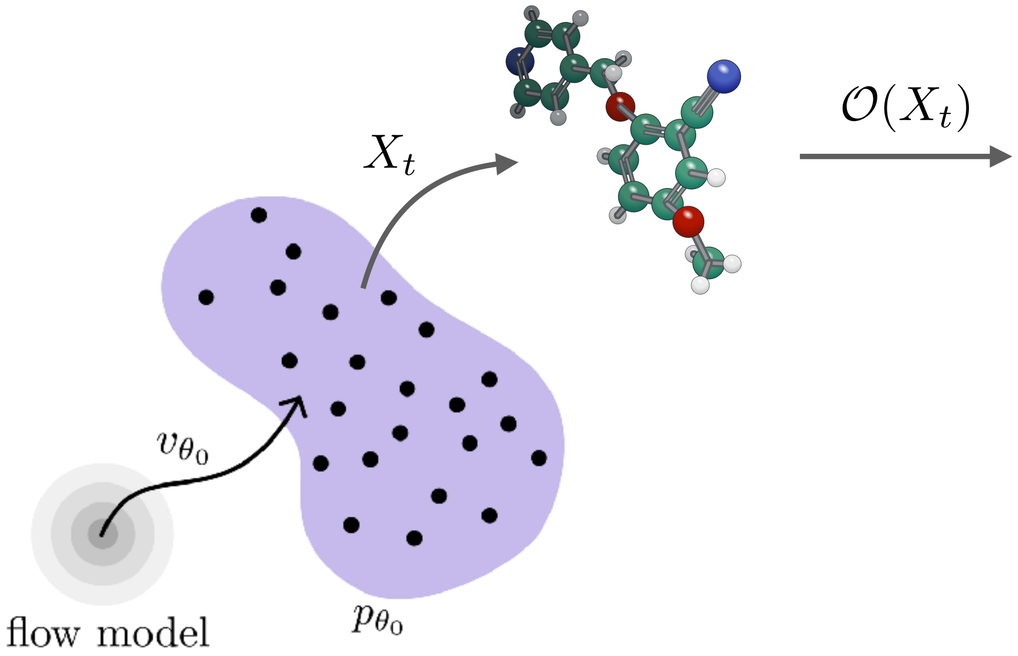
LLM as JUDGE?

Natural-language discovery goal

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

Discovery region

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



$\mathcal{O}(X_t)$

Observables

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

LLM as JUDGE?

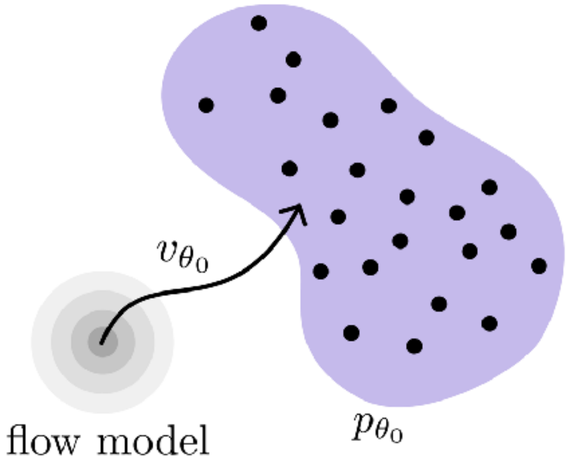
Natural-language discovery goal

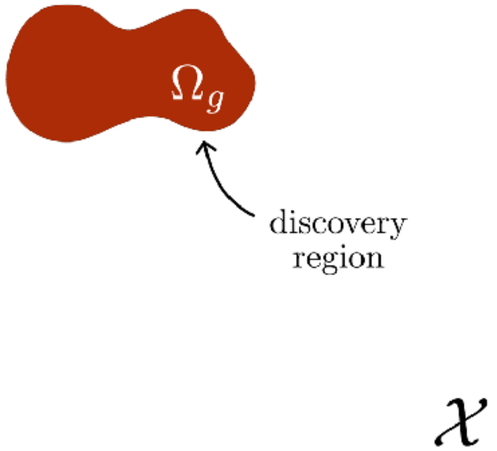
$g =$ *'generate peptides with promising anticancer activity'*

Discovery region

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

LLM as JUDGE?





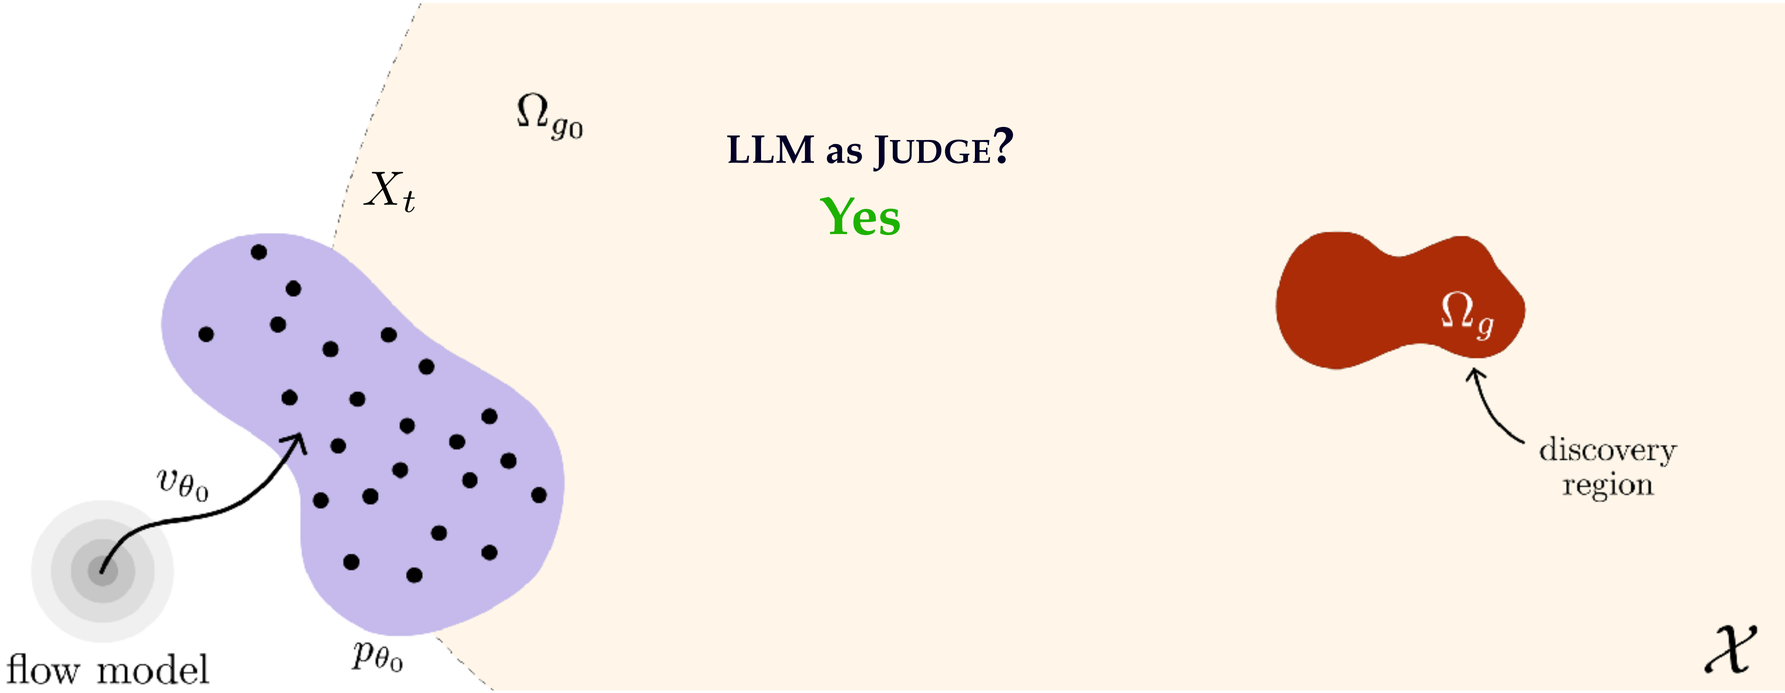
All samples X_t might be all excessively far from Ω_g

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)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

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

LLM-JUDGE:

For each $x_i \in X_t$ set $\ell_t(x_i) \in \{+, -, \varnothing\}$

$$\mathcal{P}_t = \{x_i \in X_t : \ell_t(x_i) = +\}, \quad \mathcal{N}_t = \{x_i \in X_t : \ell_t(x_i) = -\}$$

$$\mathcal{D}_t = \{(x^+, x^-) : x^+ \in \mathcal{P}_t,\ x^- \in \mathcal{N}_t\}$$

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

$D_t \leftarrow$ **LLM-JUDGE** (X_t, g_t)

$\theta_{t+1} \leftarrow$ **DPO** $(\mathcal{L}_t(\theta_t, \mathcal{D}_t))$

Return μ_{θ_T}

LLM-JUDGE:

For each $x_i \in X_t$ set $\ell_t(x_i) \in \{+, -, \varnothing\}$

$$\mathcal{P}_t = \{x_i \in X_t : \ell_t(x_i) = +\}, \quad \mathcal{N}_t = \{x_i \in X_t : \ell_t(x_i) = -\}$$

$$\mathcal{D}_t = \{(x^+, x^-) : x^+ \in \mathcal{P}_t,\ x^- \in \mathcal{N}_t\}$$

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^n \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

$D_t \leftarrow$ **LLM-JUDGE** (X_t, g_t)

$\theta_{t+1} \leftarrow$ **DPO** $(\mathcal{L}_t(\theta_t, \mathcal{D}_t))$

Return μ_{θ_T}

LLM-JUDGE:

For each $x_i \in X_t$ set $\ell_t(x_i) \in \{+, -, \varnothing\}$

$$\mathcal{P}_t = \{x_i \in X_t : \ell_t(x_i) = +\}, \quad \mathcal{N}_t = \{x_i \in X_t : \ell_t(x_i) = -\}$$

$$\mathcal{D}_t = \{(x^+, x^-) : x^+ \in \mathcal{P}_t,\ x^- \in \mathcal{N}_t\}$$

Proposition: each LFD round is a KL-regularized step toward an **implicit reward**.

$$p_t^\star \propto p_{\theta_t} e^{f_{g_t}/\beta} \quad \Longleftrightarrow \quad p_t^\star = \arg\max_q \{\mathbb{E}_q[f_{g_t}] - \beta \mathcal{D}_{\mathrm{KL}}(q \,\|\, p_{\theta_t})\}$$

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

$D_t \leftarrow$ **LLM-JUDGE** (X_t, g_t)

$\theta_{t+1} \leftarrow$ **DPO** $(\mathcal{L}_t(\theta_t, \mathcal{D}_t))$

Return μ_{θ_T}

LLM-JUDGE:

For each $x_i \in X_t$ set $\ell_t(x_i) \in \{+, -, \varnothing\}$

$$\mathcal{P}_t = \{x_i \in X_t : \ell_t(x_i) = +\}, \quad \mathcal{N}_t = \{x_i \in X_t : \ell_t(x_i) = -\}$$

$$\mathcal{D}_t = \{(x^+, x^-) : x^+ \in \mathcal{P}_t,\ x^- \in \mathcal{N}_t\}$$

LFD $\approx$ Iterative DPO on implicit LLM rewards

LLM-to-Flow Distillation (LFD)

Require: μ_{θ_0}, goal g, observables $\mathcal{O}$

$g_{-1} \leftarrow g$

For $t = 0, \ldots, T-1$

$X_t = \{x_i\}_{i=1}^{n} \sim p_{\theta_t}$ observe $\mathcal{O}(X_t)$

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

$D_t \leftarrow$ **LLM-JUDGE** (X_t, g_t)

$\theta_{t+1} \leftarrow$ **DPO** $(\mathcal{L}_t(\theta_t, \mathcal{D}_t))$

Return μ_{θ_T}

LLM-JUDGE:

For each $x_i \in X_t$ set $\ell_t(x_i) \in \{+, -, \varnothing\}$

$$\mathcal{P}_t = \{x_i \in X_t : \ell_t(x_i) = +\}, \quad \mathcal{N}_t = \{x_i \in X_t : \ell_t(x_i) = -\}$$

$$\mathcal{D}_t = \{(x^+, x^-) : x^+ \in \mathcal{P}_t,\ x^- \in \mathcal{N}_t\}$$

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

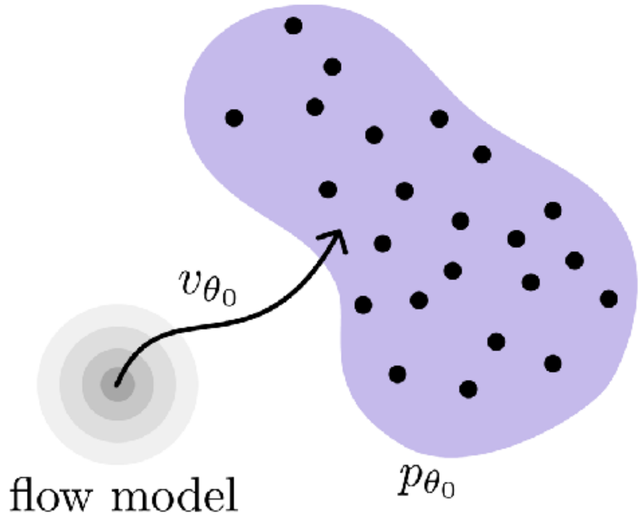
From Rewards to Natural-Language Goals

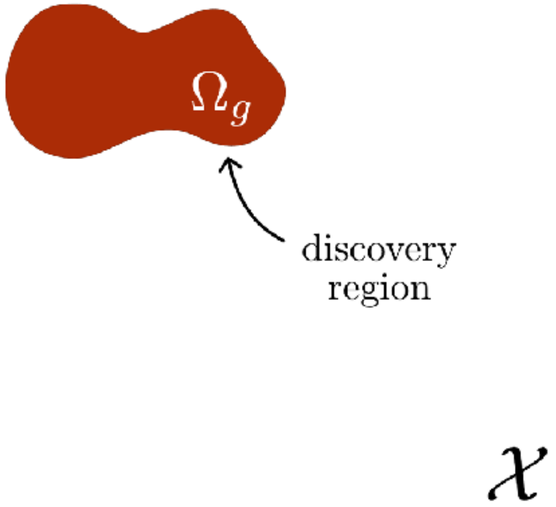
Natural-language discovery goal

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

Discovery region

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





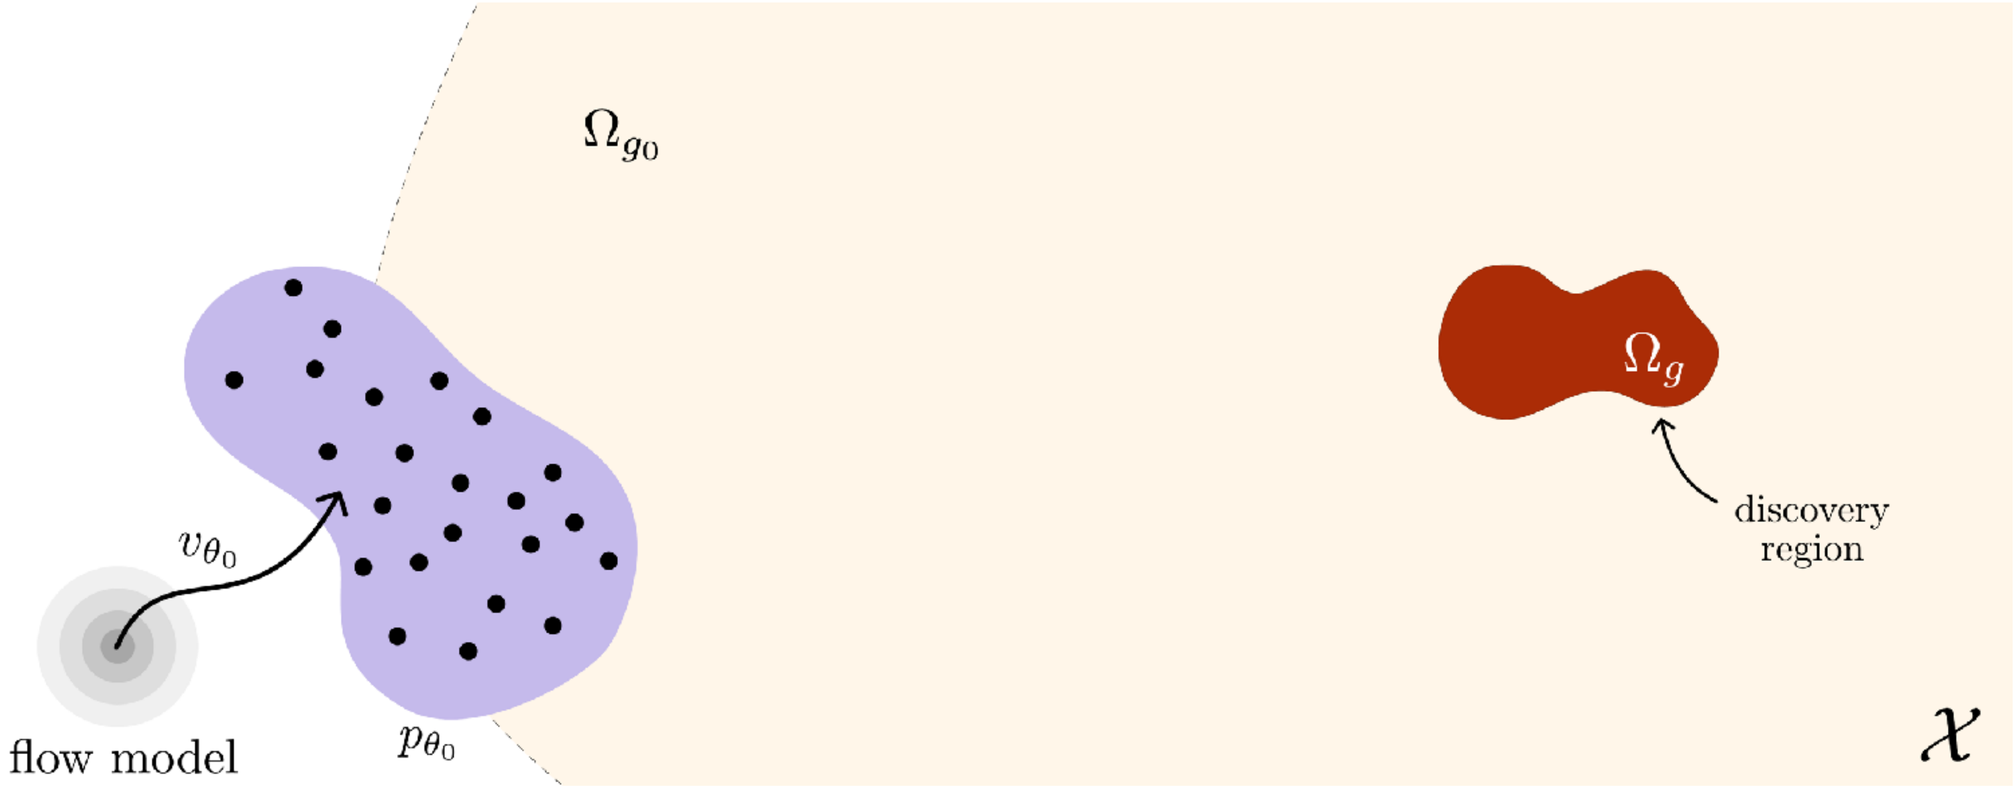
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



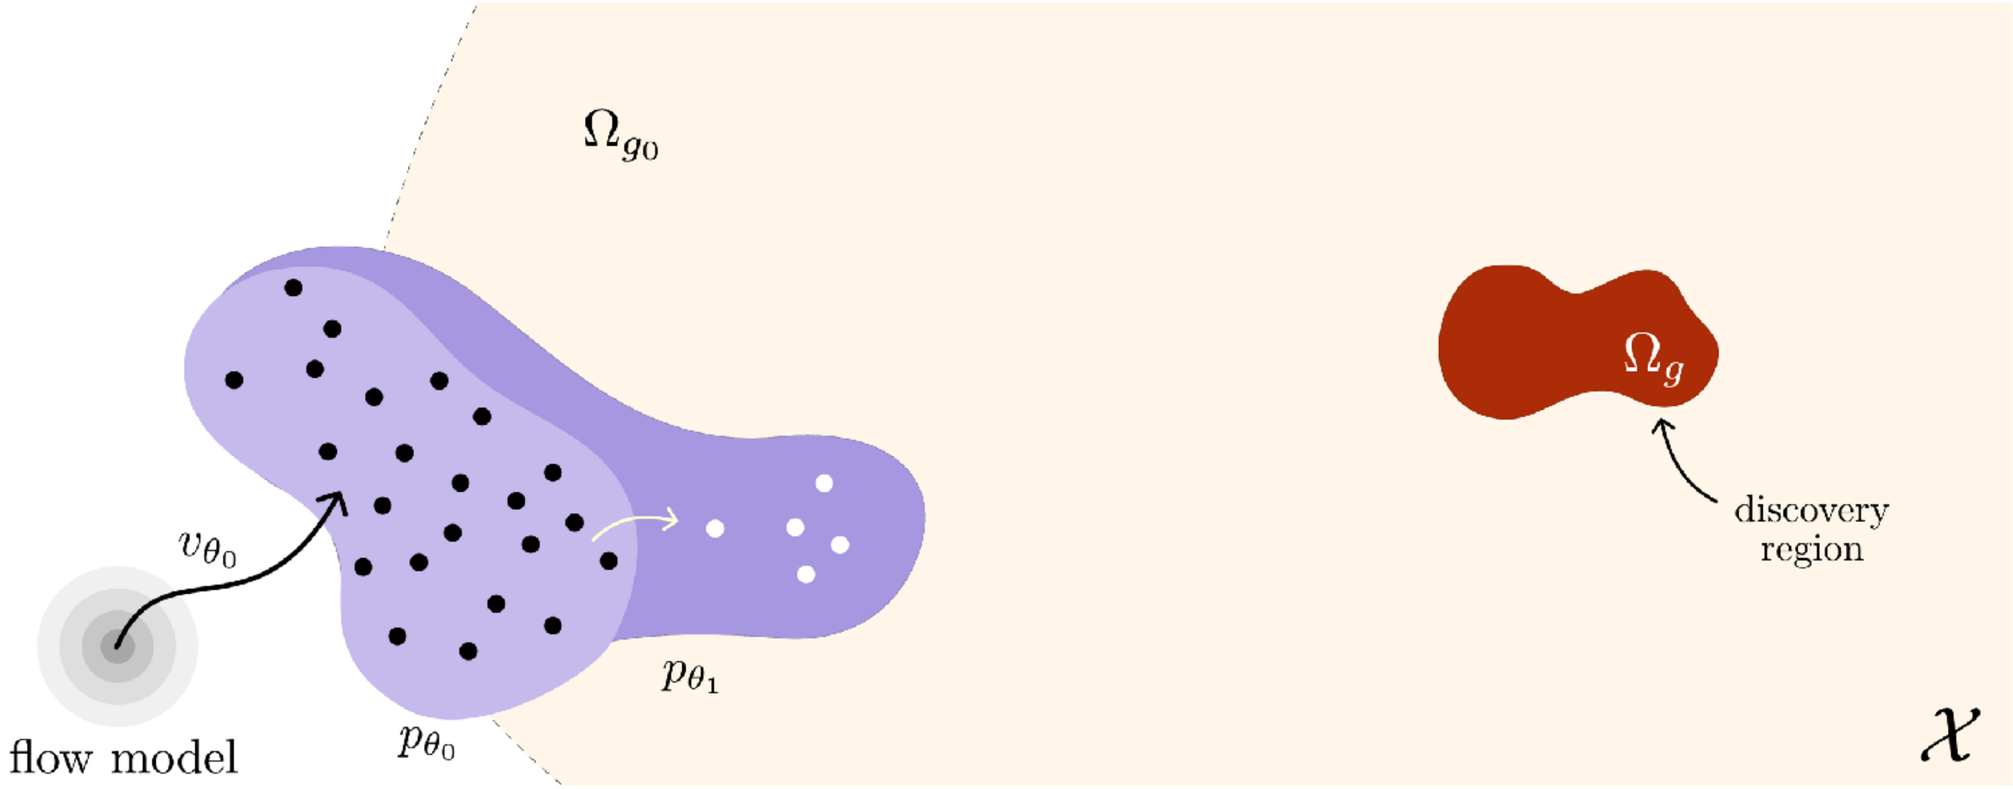
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



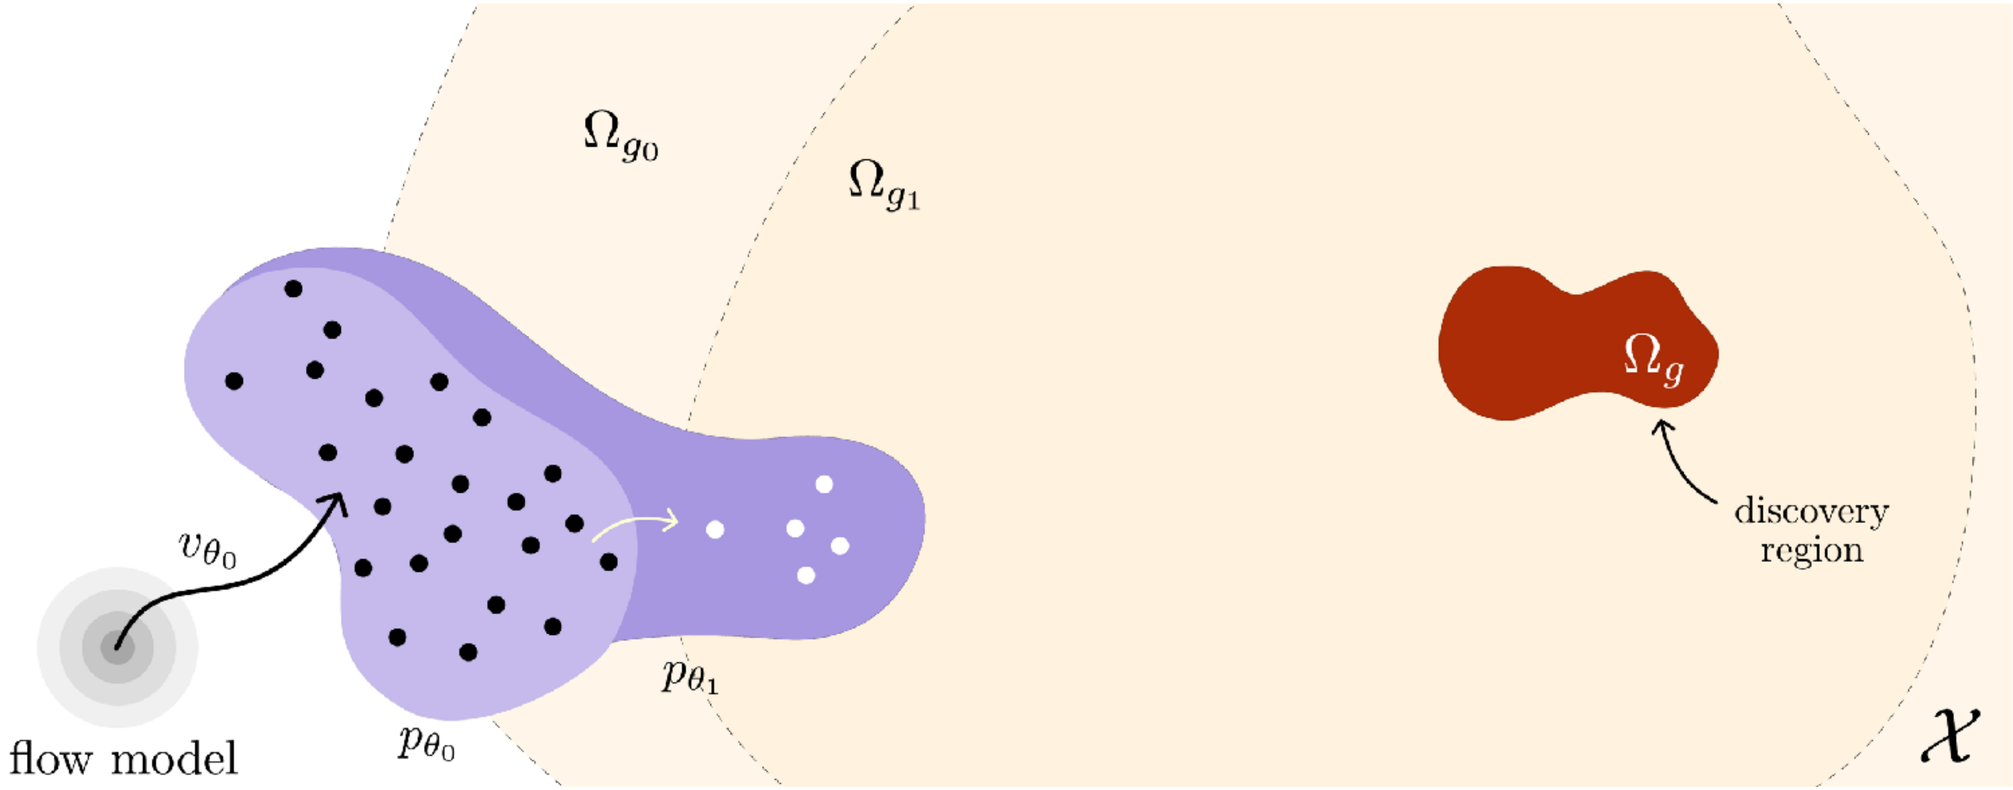
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



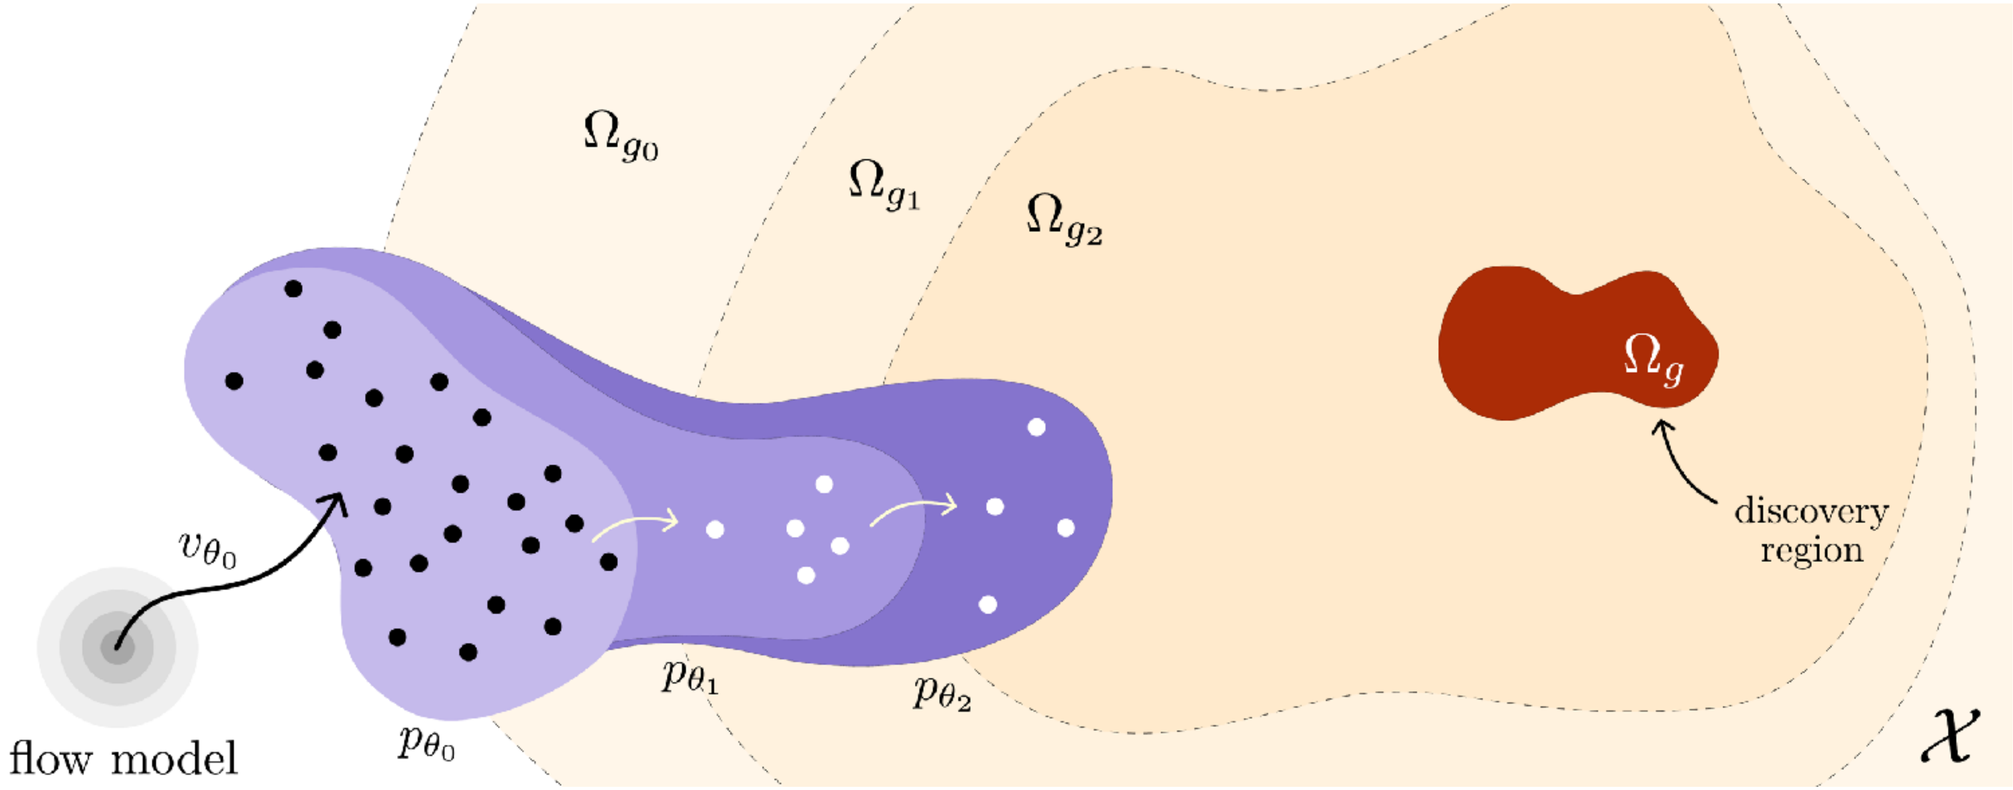
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



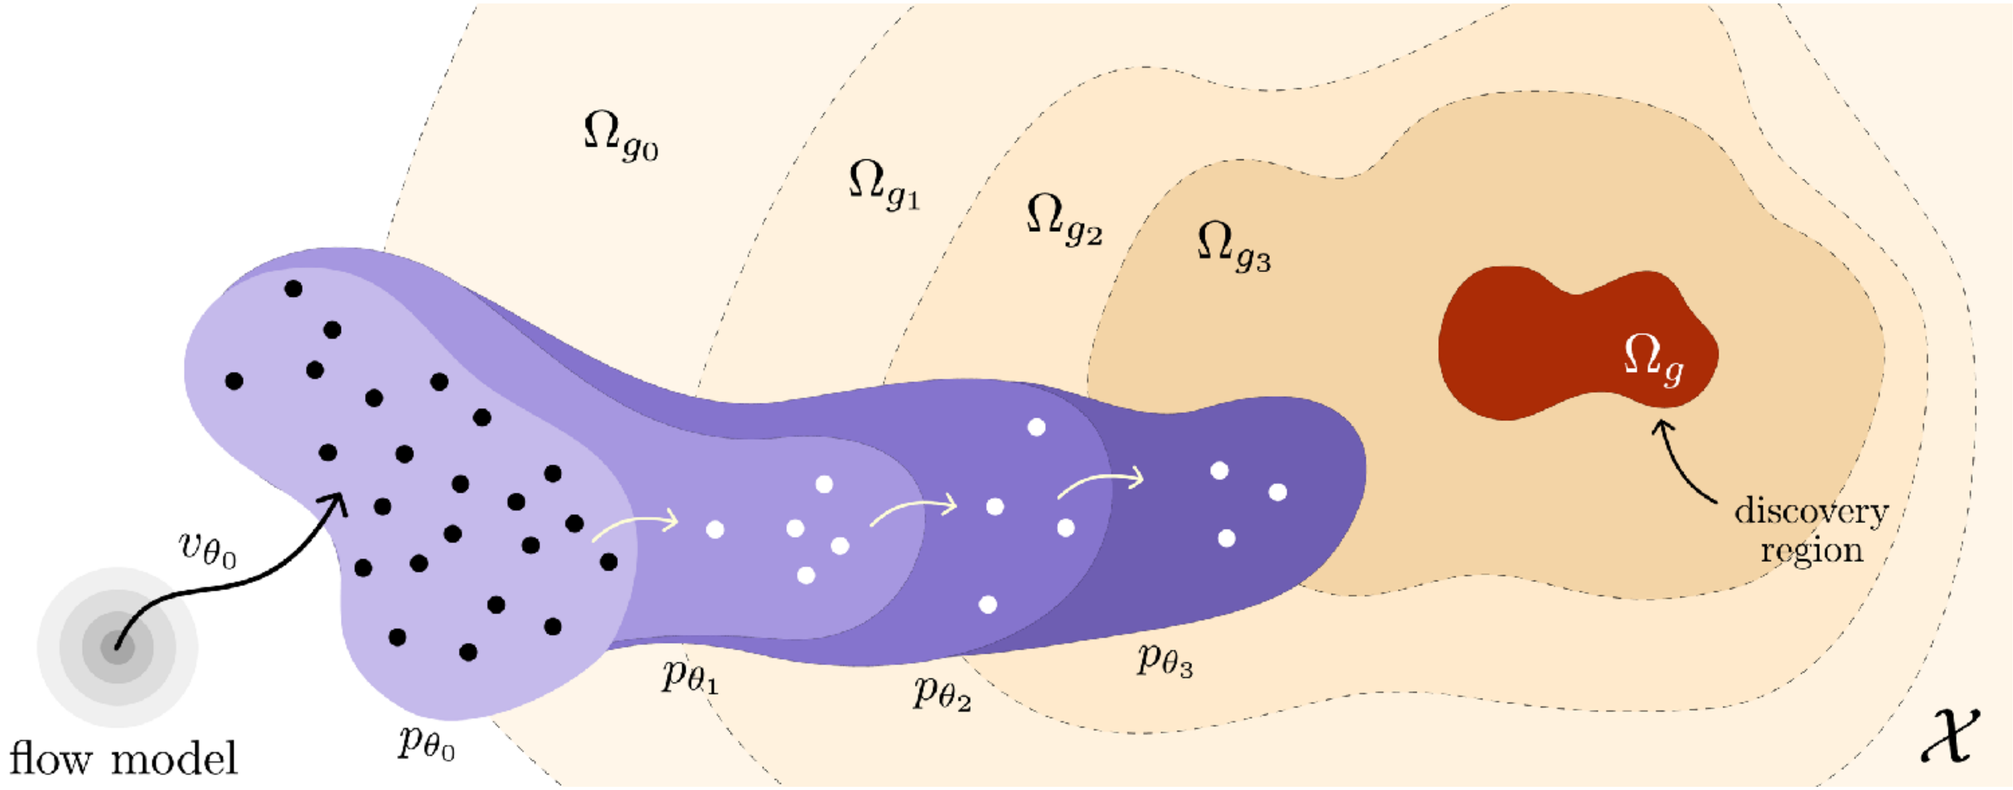
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



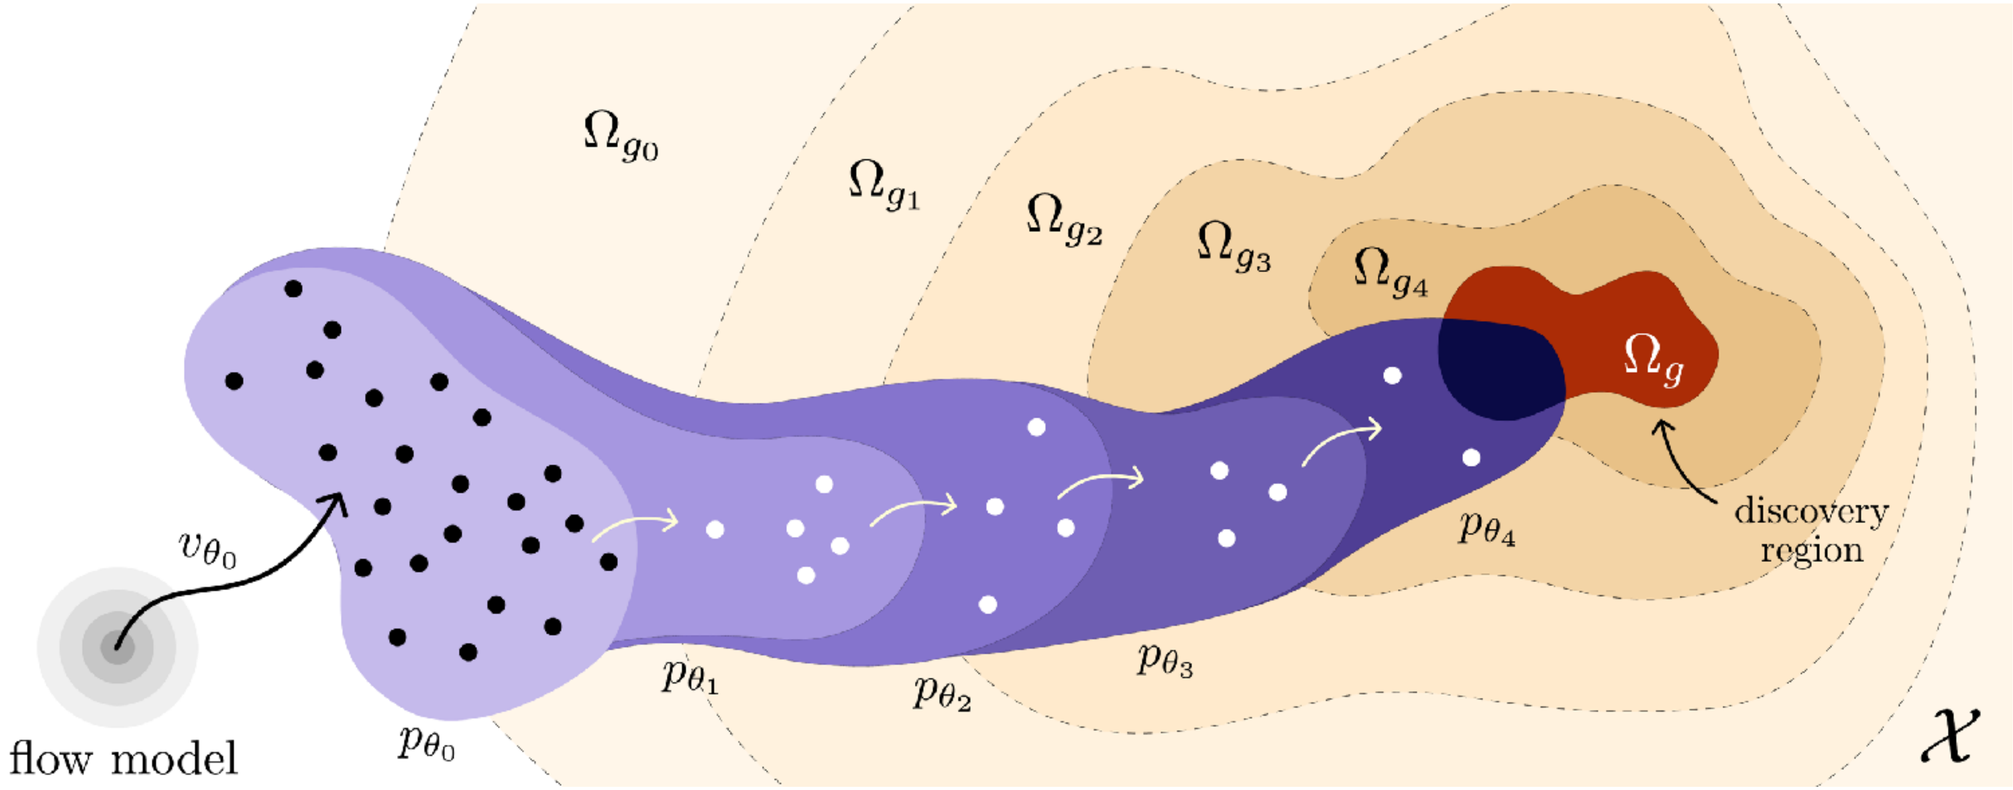
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



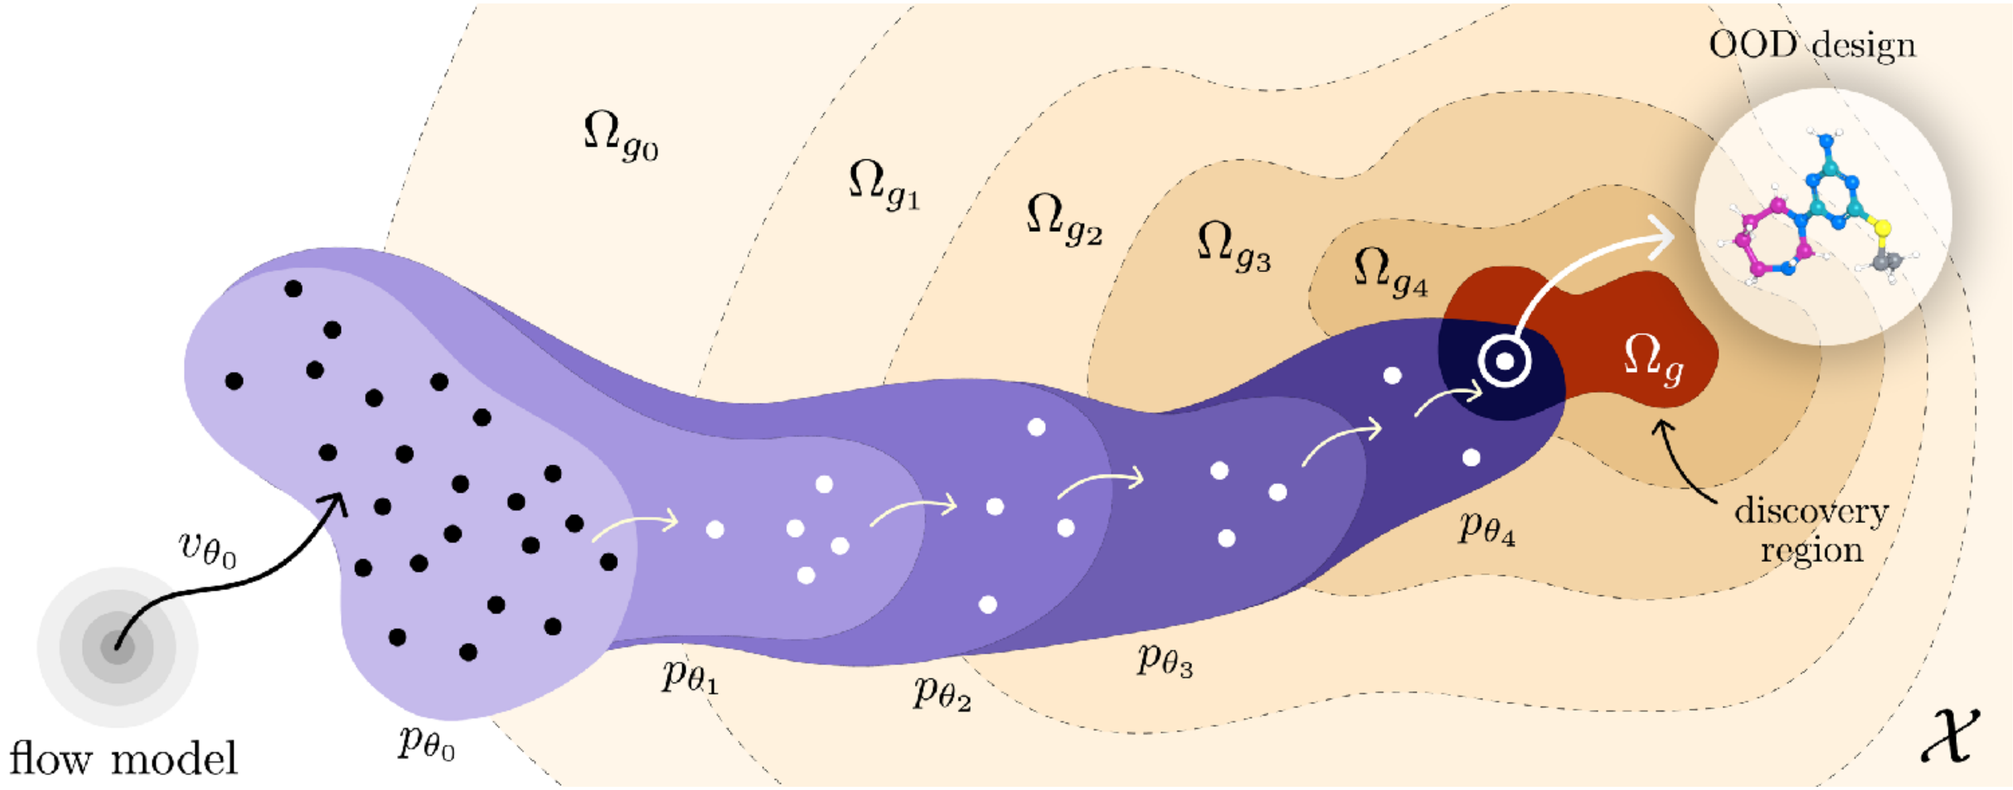
From Rewards to Natural-Language Goals

Natural-language discovery goal

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

Discovery region

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



LFD for Peptide Design

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

LFD for Peptide Design

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

GOAL 2: Generate peptides with promising **mammalian-cell-penetrating activity** while **minimizing toxicity** to mammalian cells.

LFD does not observe any proxy for anticancer or cell-penetrating activity!

LFD for Peptide Design

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

GOAL 2: Generate peptides with promising **mammalian-cell-penetrating activity** while **minimizing toxicity** to mammalian cells.

LFD does not observe any proxy for anticancer or cell-penetrating activity!

	Therapeutic Peptides (CPP)			Therapeutic Peptides (Anti-Cancer)		
Method	**pLM4CPPs (%) ↑**	**HemoPI2 (%) ↓**	**Distance from prior ↑**	**AntiCP2-alt. (%) ↑**	**HemoPI2 (%) ↓**	**Distance from prior ↑**
Pre-trained	8.81 ± 3.20	40.84 ± 1.96	0.06 ± 0.04	26.43 ± 2.91	**40.85** ± 1.89	0.00 ± 0.00
DPO	10.39 ± 3.19	40.49 ± 1.92	0.07 ± 0.04	25.29 ± 3.19	42.77 ± 1.81	0.08 ± 0.04
DIFFUSIONNFT	57.04 ± 6.22	38.87 ± 1.62	0.29 ± 0.03	33.26 ± 2.76	45.61 ± 2.39	0.16 ± 0.03
GUIDANCE	8.49 ± 3.30	43.45 ± 1.85	0.06 ± 0.04	35.94 ± 2.33	47.25 ± 1.81	0.29 ± 0.03
LFD	**69.06** ± 6.36	**37.02** ± 1.69	**0.34** ± 0.03	**53.54** ± 3.32	42.48 ± 2.04	**0.31** ± 0.03

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

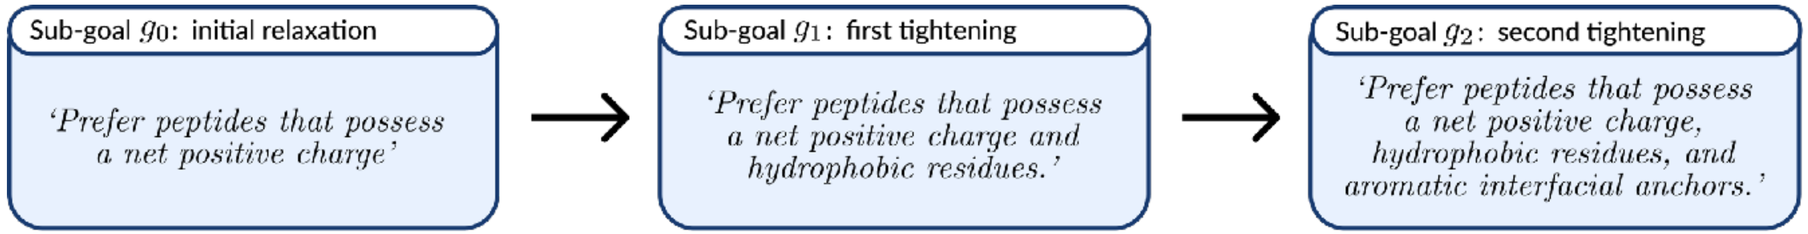
INSIGHT: LFD can align domain-specific scientific generators to high-level, natural-language discovery goals, **without utilizing any proxy reward.**

LFD for Peptide Design: Interpretability

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

LFD for Peptide Design: Interpretability

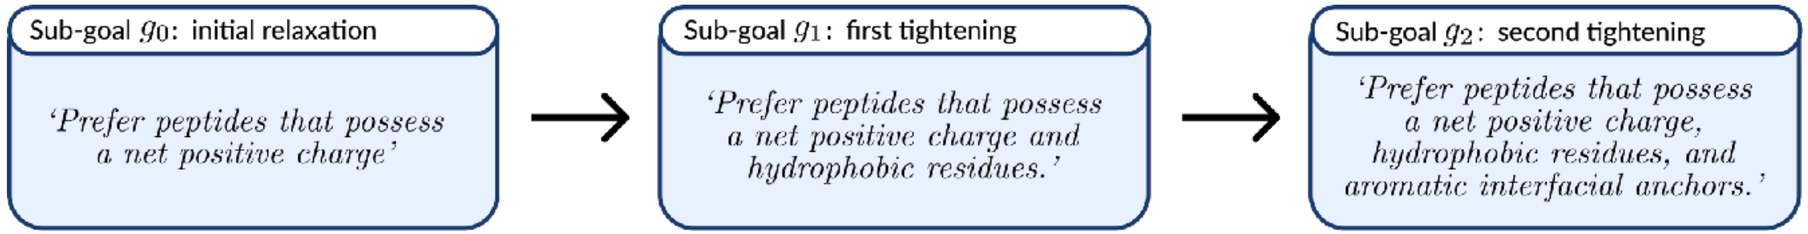
GOAL 1: Generate peptides with promising **anticancer activity** while maintaining **low toxicity** toward healthy cells.



LFD generates sub-goals interpretable by scientists.

LFD for Peptide Design: Interpretability

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.

LFD Surpasses Reward-Adaptation in Molecular Design Tasks

> **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 Surpasses Reward-Adaptation in Molecular Design Tasks

> **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**.

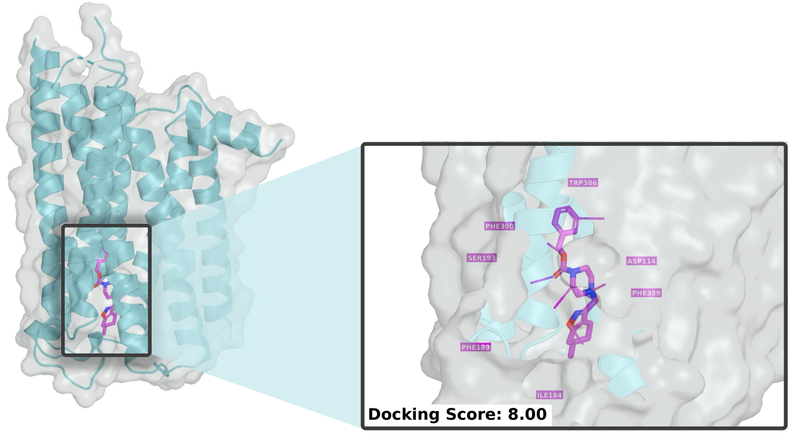
LFD does not observe any docking proxy!

LFD Surpasses Reward-Adaptation in Molecular Design Tasks

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**.

	Small Molecules (D2 Receptor)		
Method	**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 reward-guided baselines on diverse **held-out docking scores.**



LFD-generated molecule docks to **D2 dopamine receptor**

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

LFD for Molecular Design: Complex Constraints + OOD Generation

> **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**.

LFD for Molecular Design: Complex Constraints + OOD Generation

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 for Molecular Design: Complex Constraints + OOD Generation

> **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**

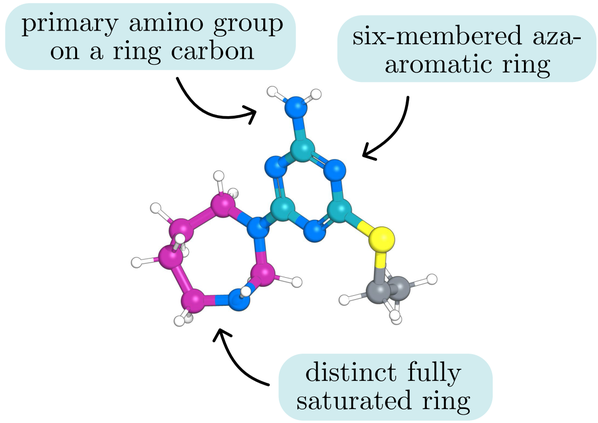
LFD for Molecular Design: Complex Constraints + OOD Generation

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.



INSIGHT: Through its adaptive LLM-generated sub-goal curriculum, LFD can adapt flow models to OOD discovery goals where state-of-the-art reward-guided methods might fail.

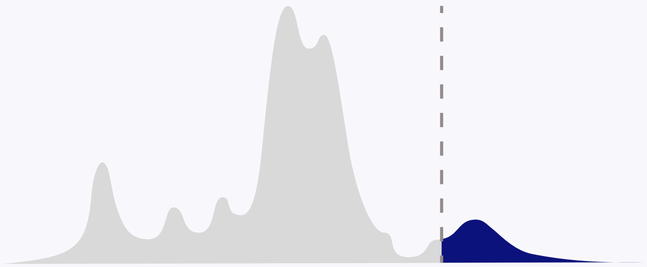
Generative Discovery

Generation of **rare and diverse structures,** potentially **far beyond available data** and **sample-efficiently**

Part I

Rare and Diverse Generation

Which candidates matter?

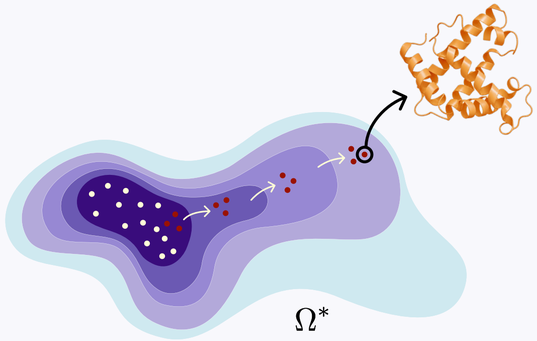


→ FDC

Part II

Out-of-Distribution Flow Modeling

Which regions can we reach?

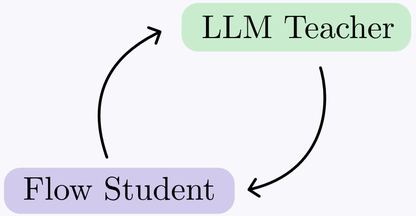


→ ActFlow

Part III

Sample-Efficient Discovery via LLM Feedback

What feedback guides us there?



→ LFD

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

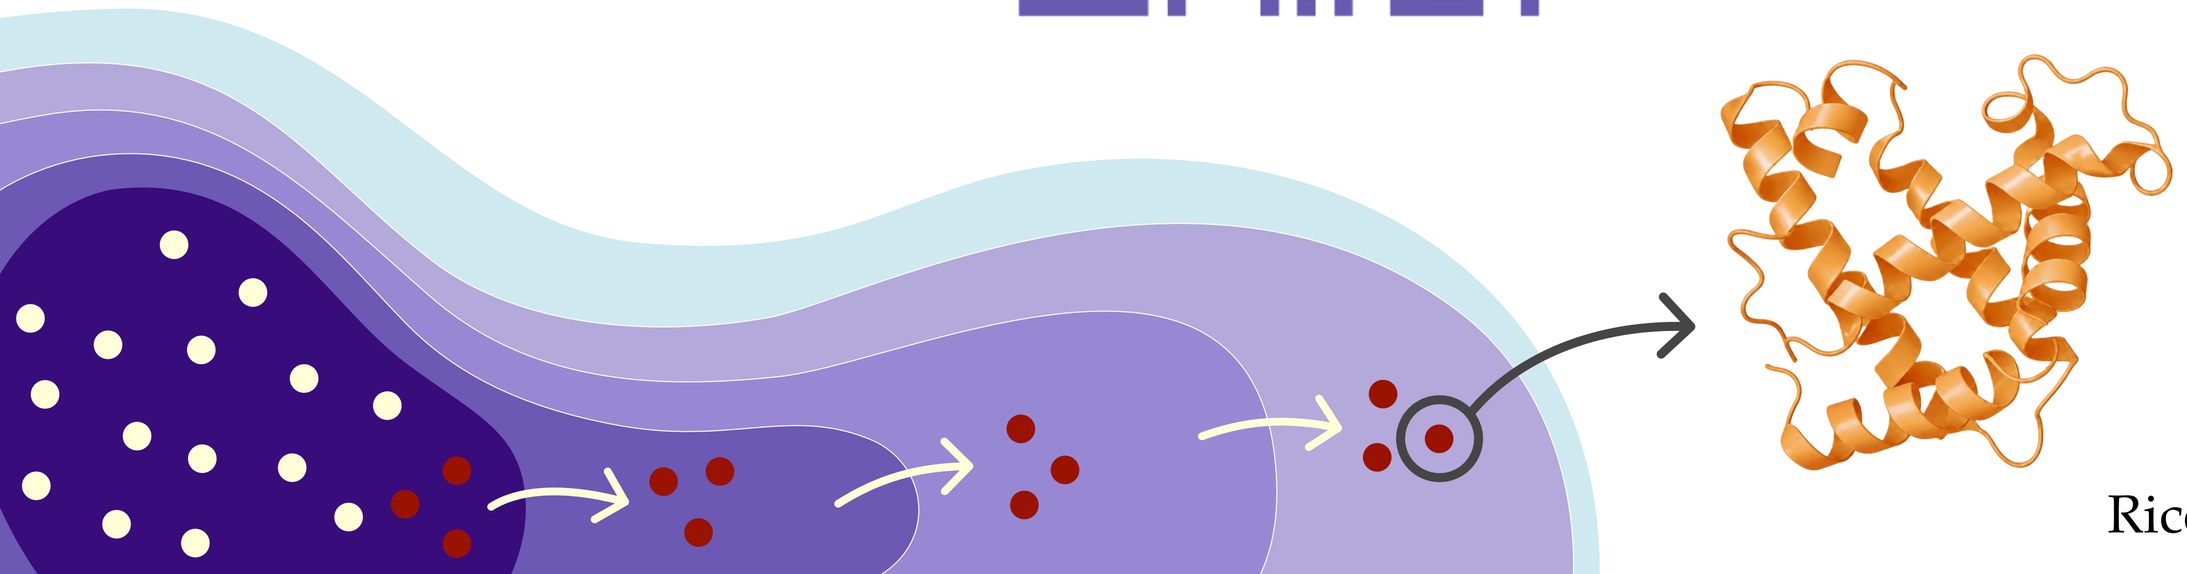
Thanks!

Papers + Code → riccardodesanti.com/research

Connect:

X: @r_de_santi

email: rdesanti@ethz.ch



Riccardo De Santi