Foundations of Generative Discovery Beyond the Data via Flow and Diffusion Models

Riccardo De Santi

Doctoral Fellow at ETH AI Center

@r_de_santi rdesanti@ethz.ch

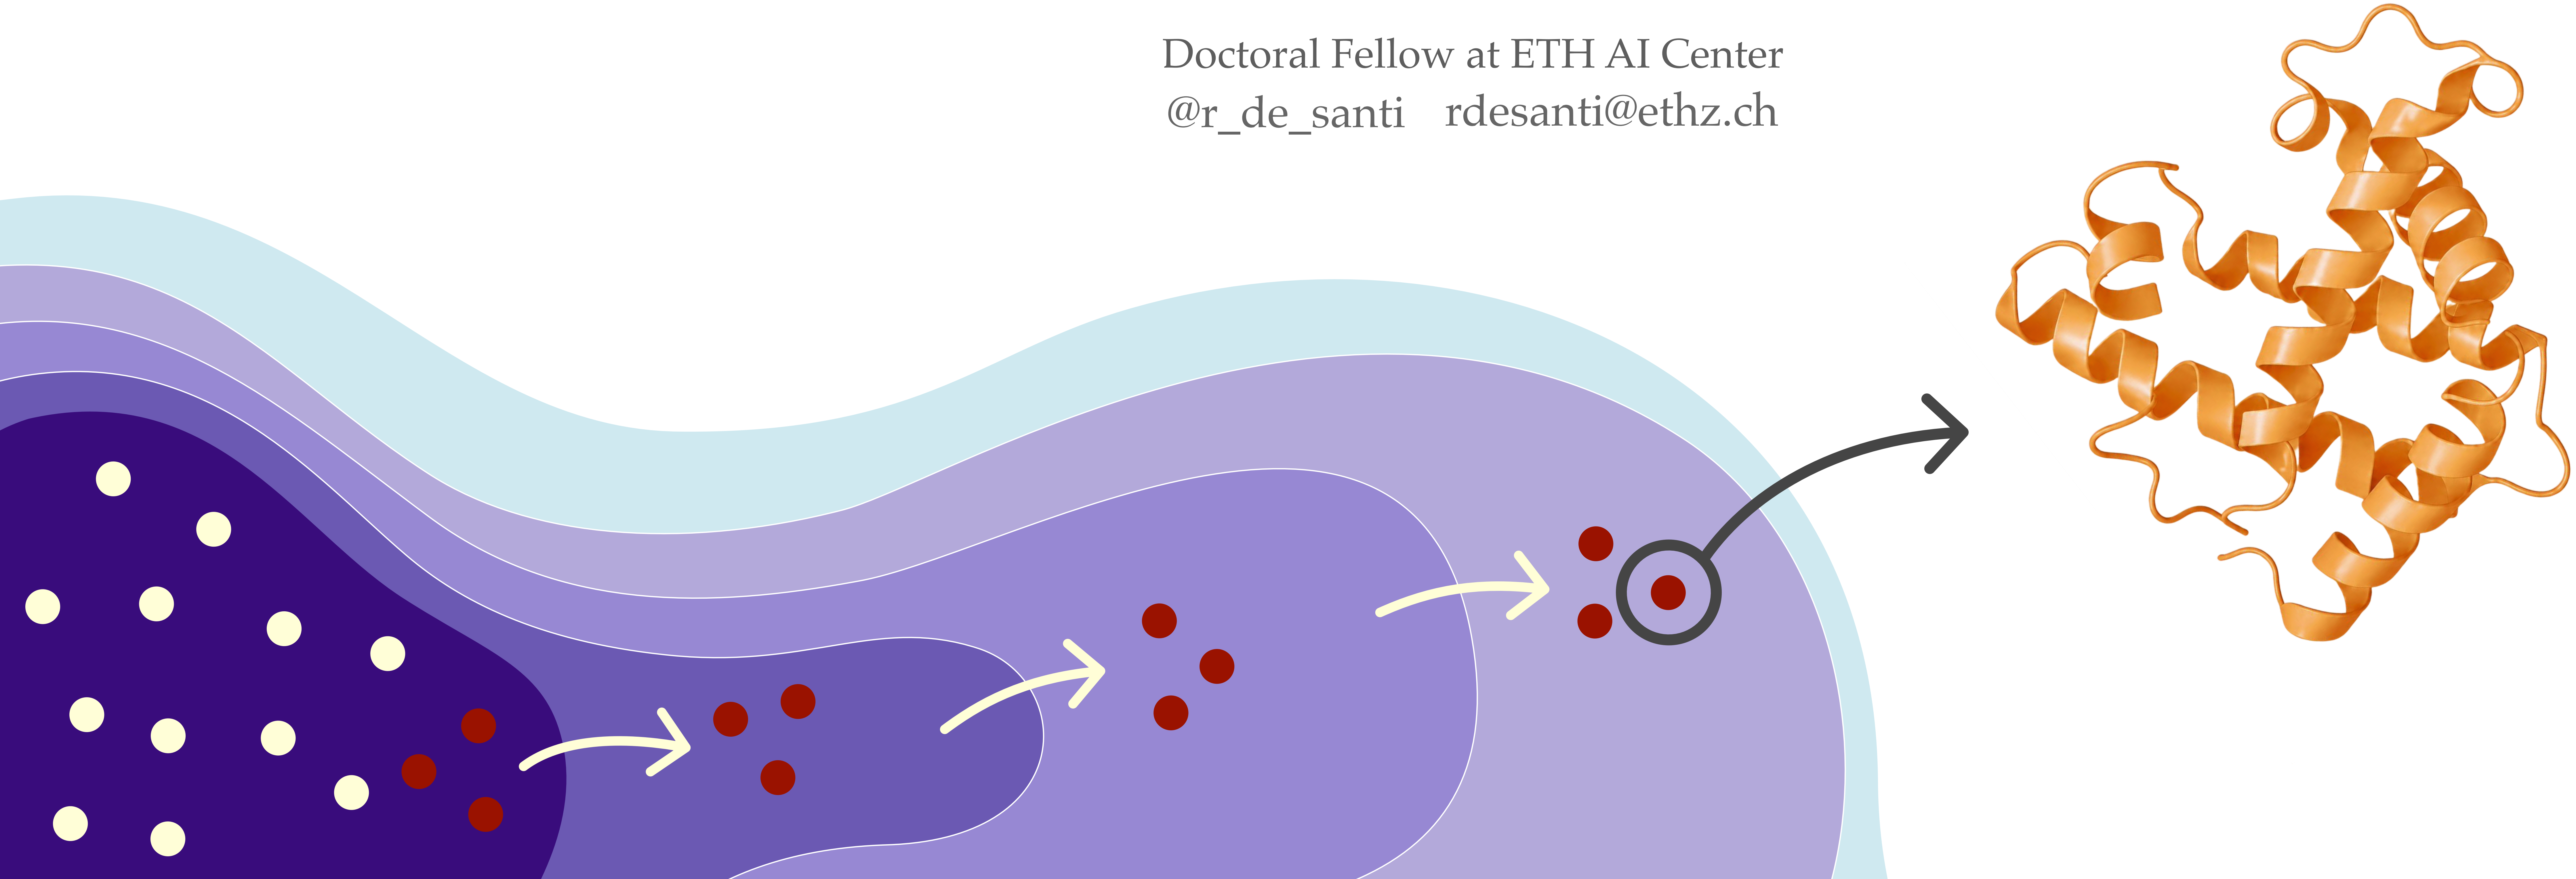

Team

Riccardo De Santi

Ya-Ping Hsieh

Kimon Protopapas

Bruce Lee

Cristian P. Jensen

Zebang Shen

Marin Vlastelica

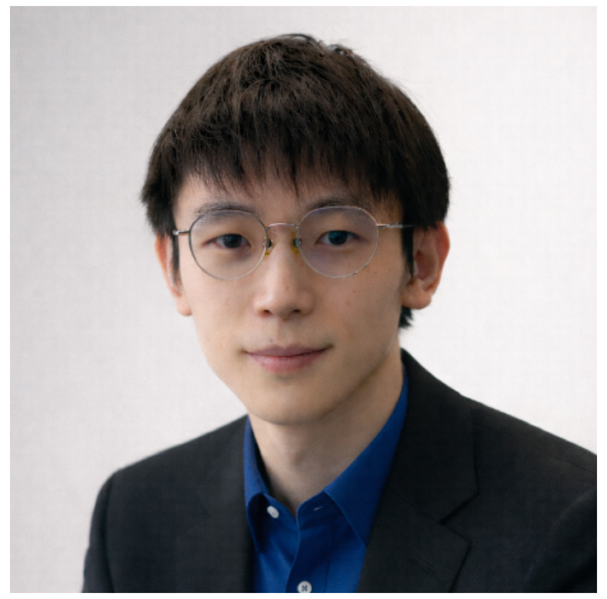

Zifan Wang

Malte Franke

Sophia Tong

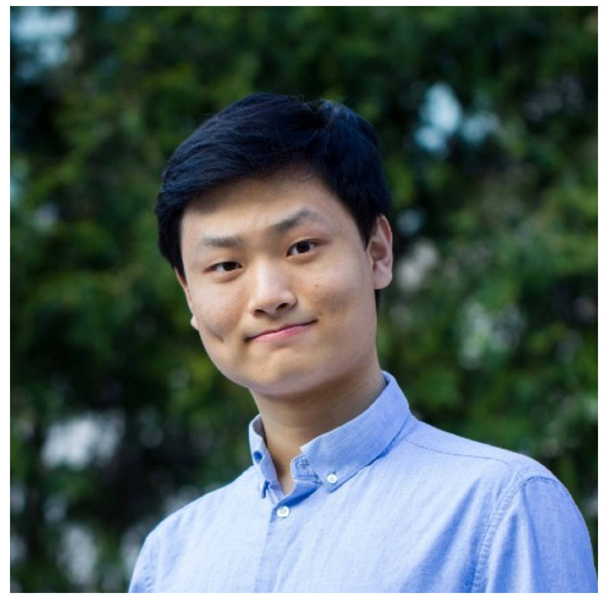

Chenghao Liu

Pranam Chatterjee

Team

Andreas Krause

Niao He

Kjell Jorner

Yisong Yue

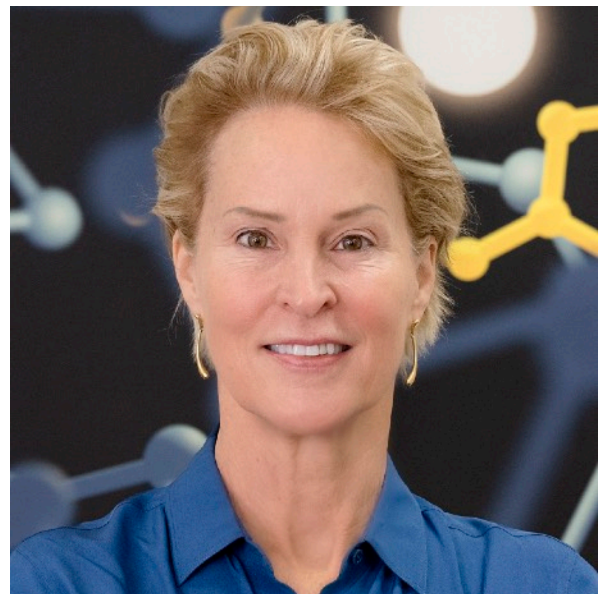

Frances H. Arnold

Generative Models

[Steven Wommack]

Generative Models for Science

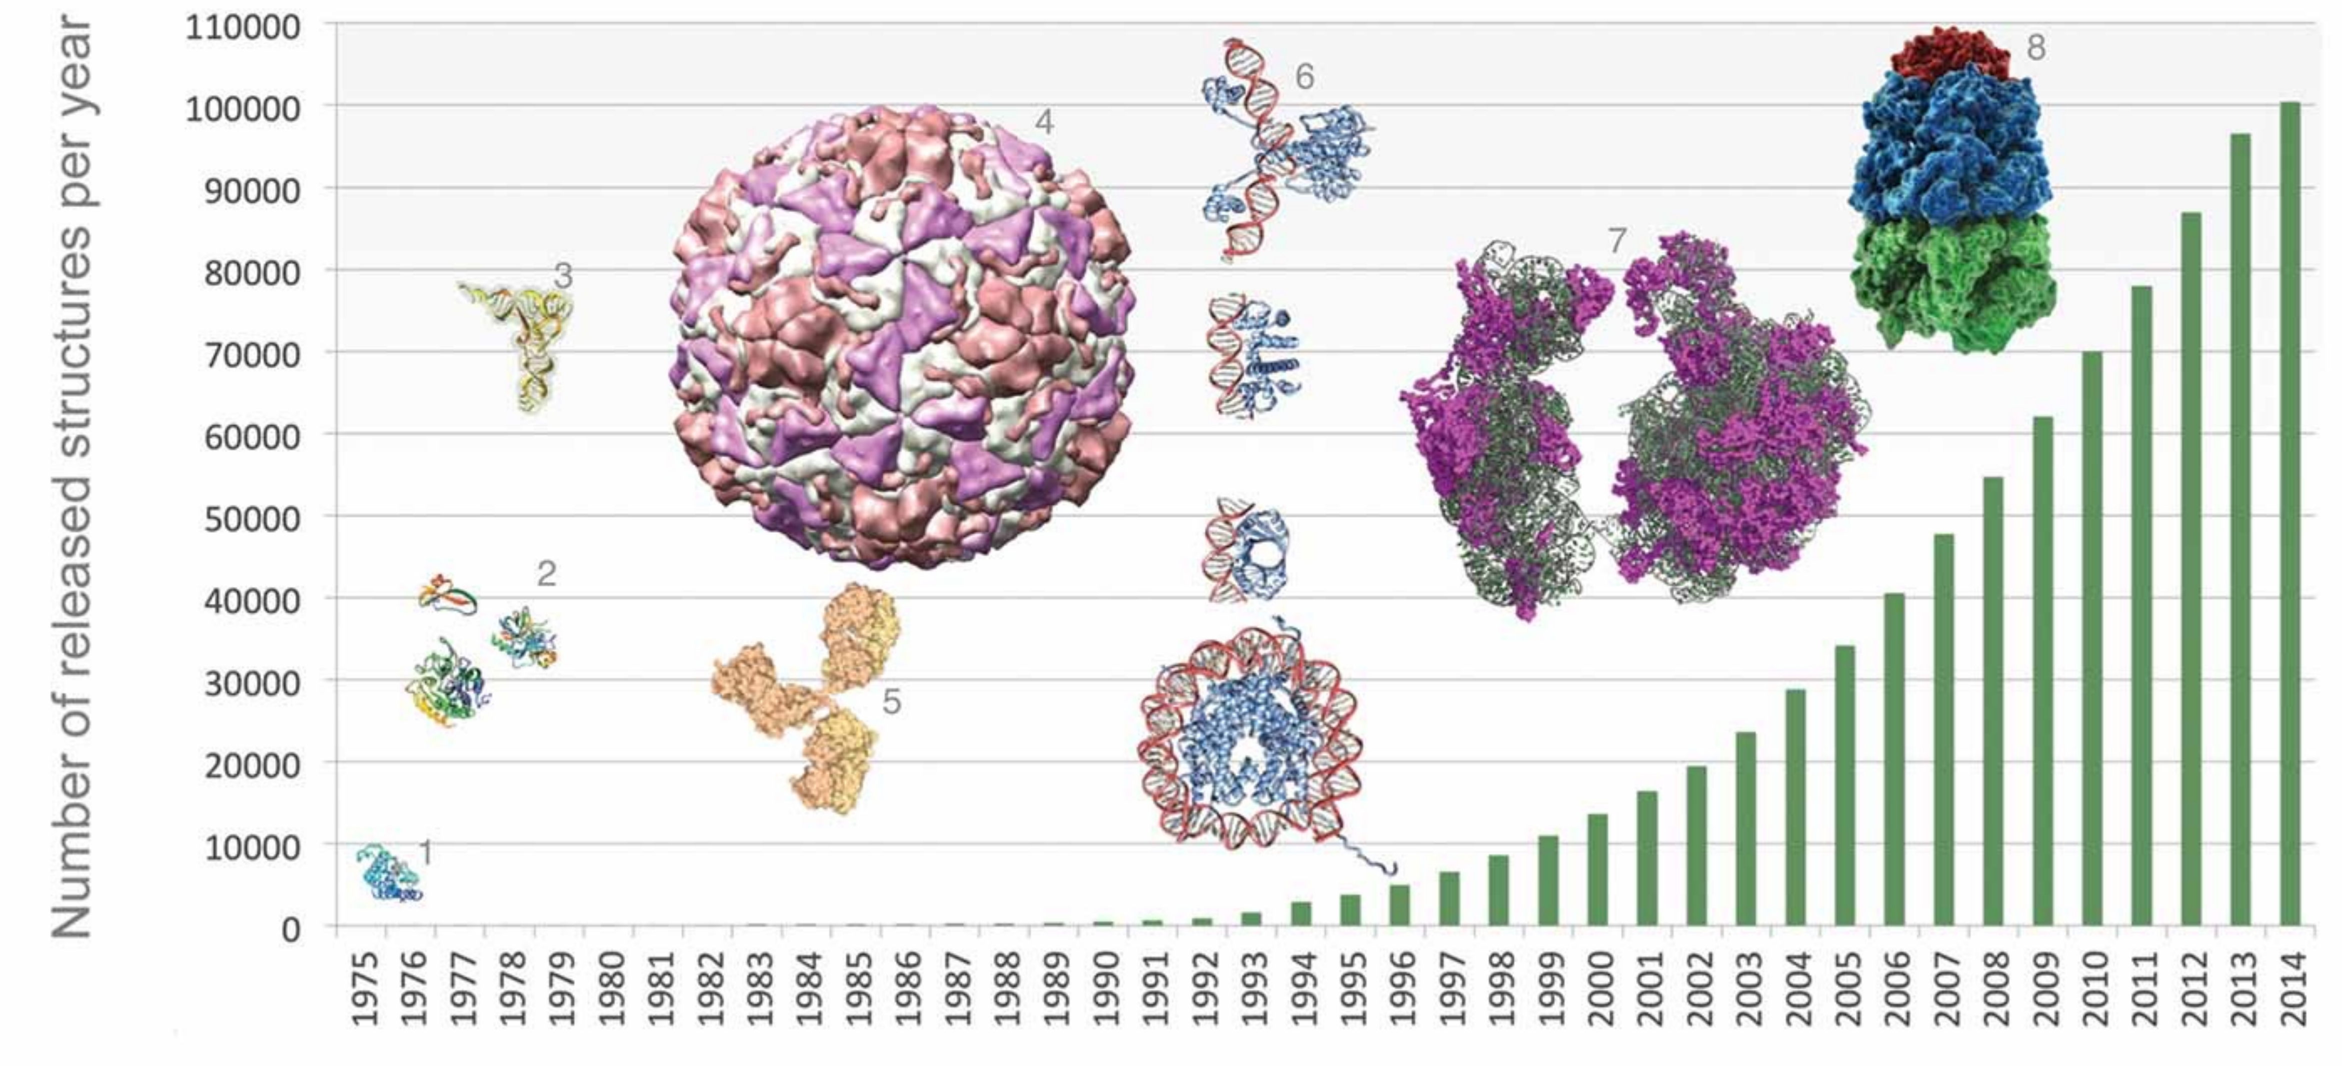

Protein Data Databank **(PDB)**

Generative Models for Science

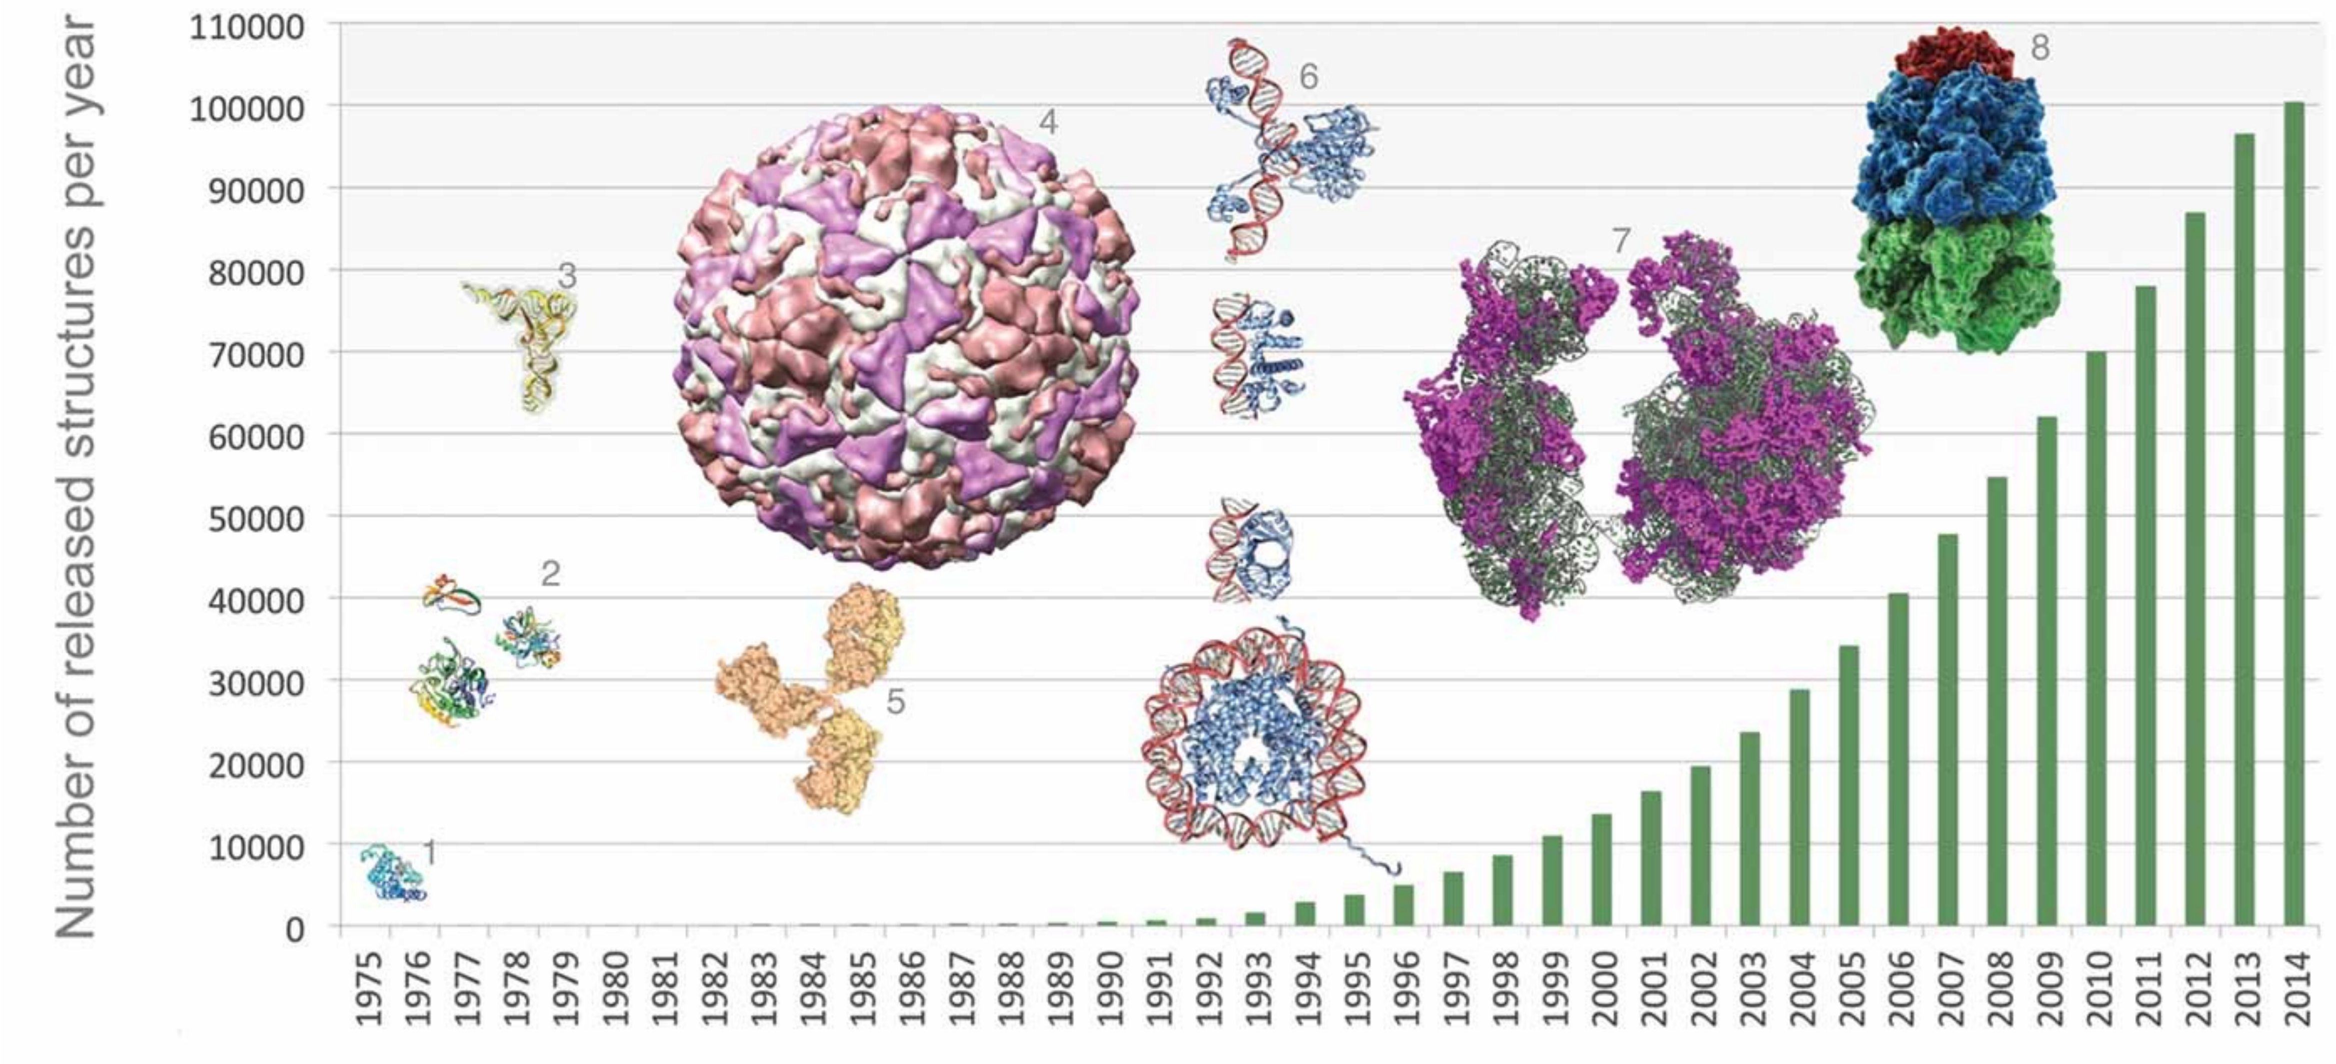

Protein Data Databank **(PDB)**

2024 Chemistry Nobel Prize

Protein structure prediction and design

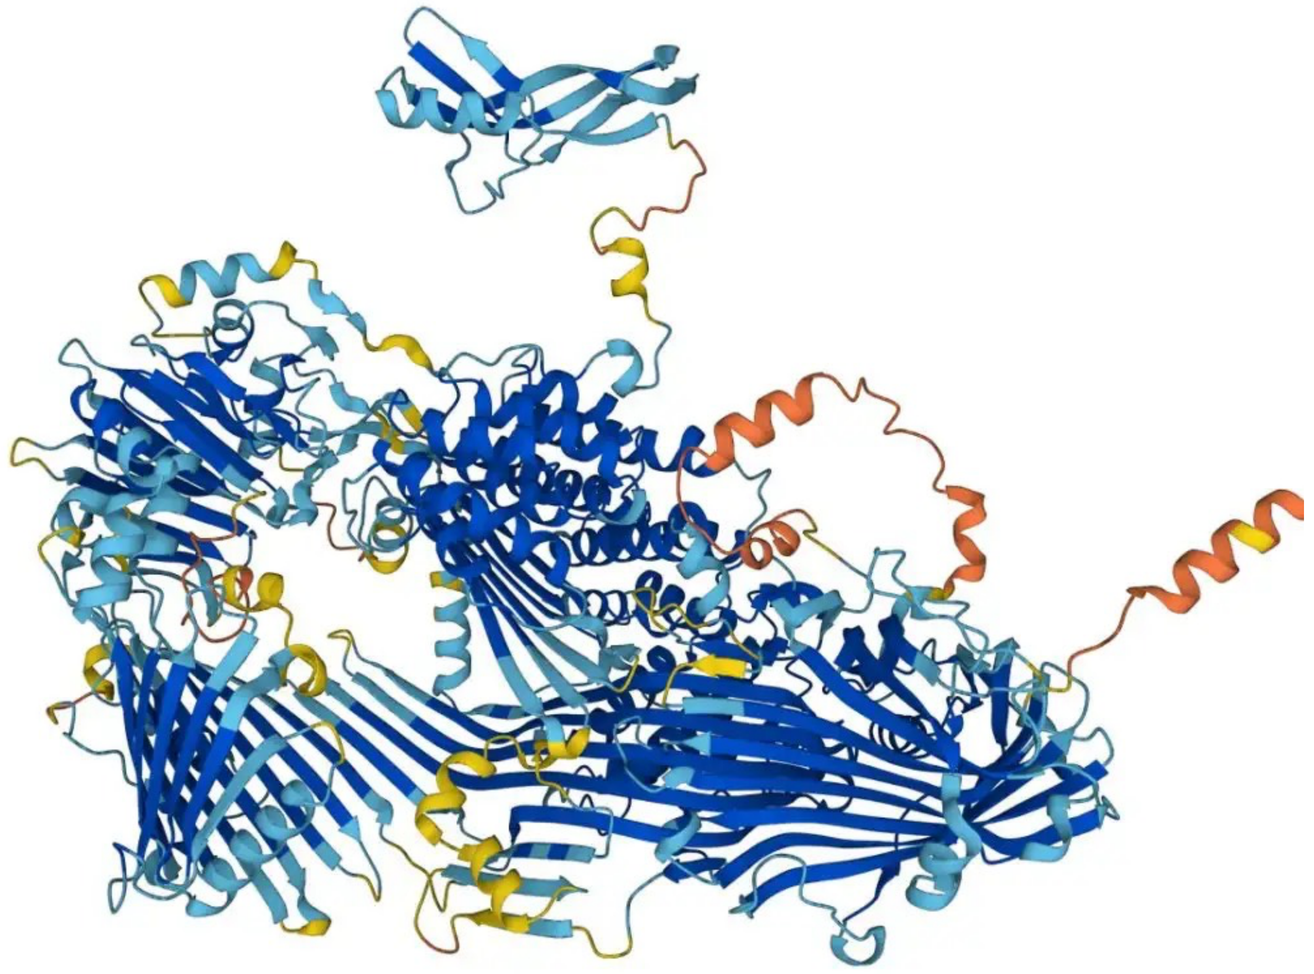

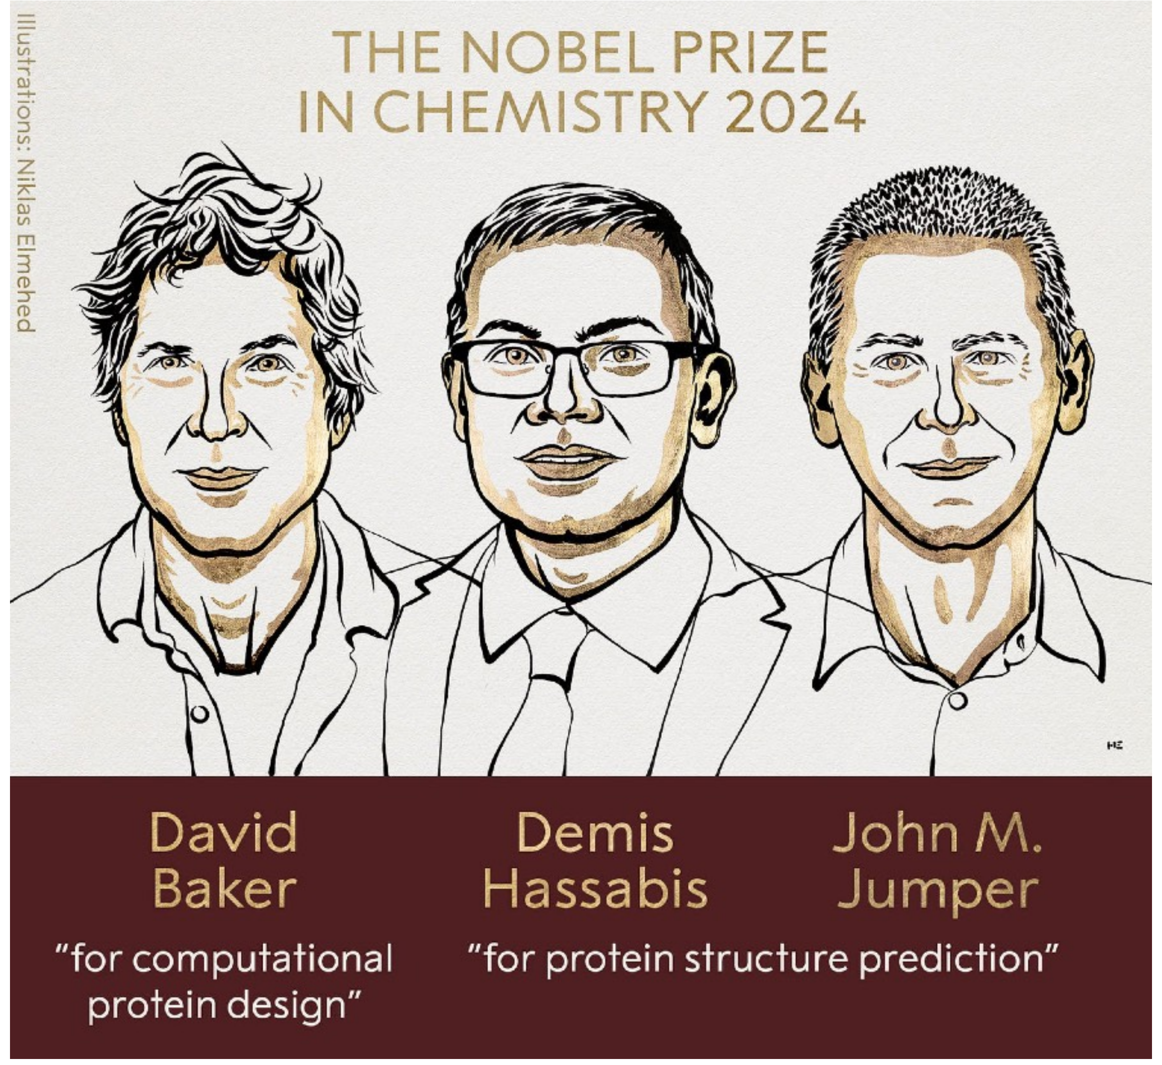

Generative ~~Modeling~~ Discovery?

What is Discovery?

black-box function (reward)

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

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

(or **action space**)

What is Discovery?

black-box function (reward)

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

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

(or **action space**)

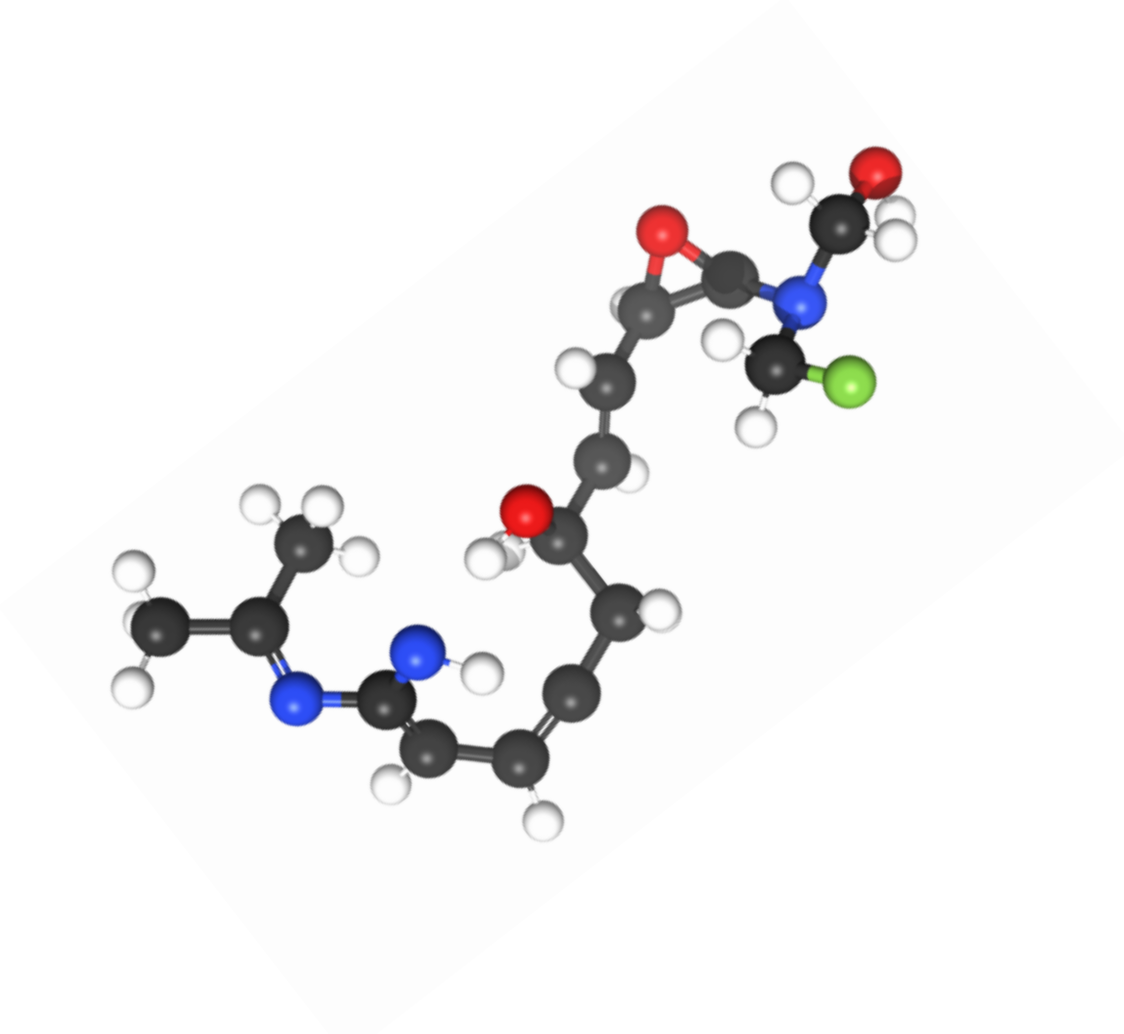

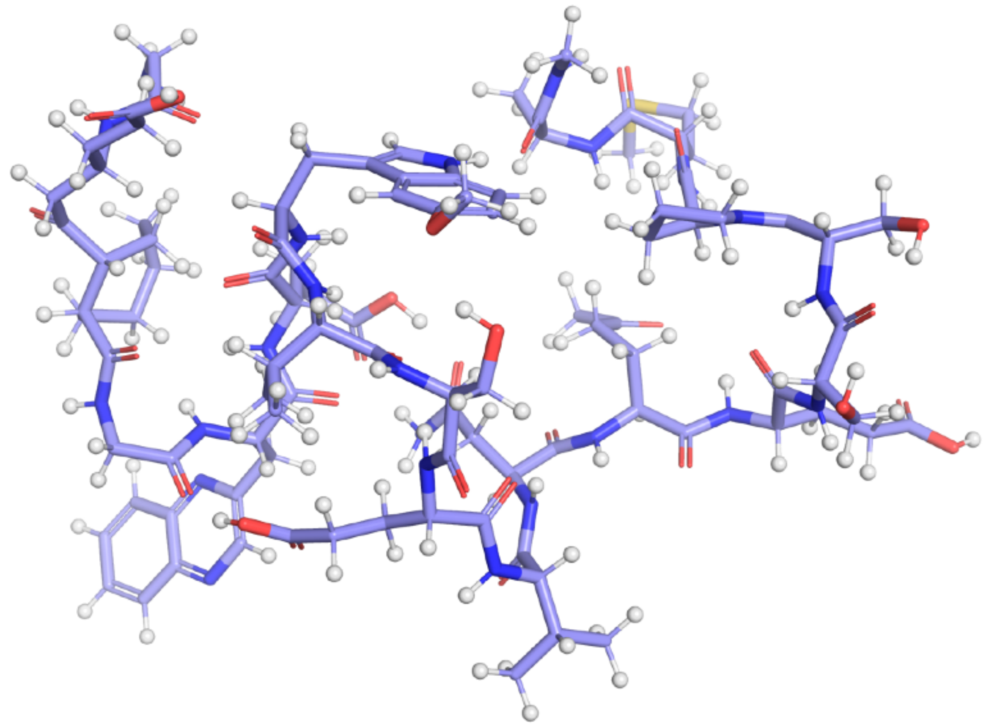

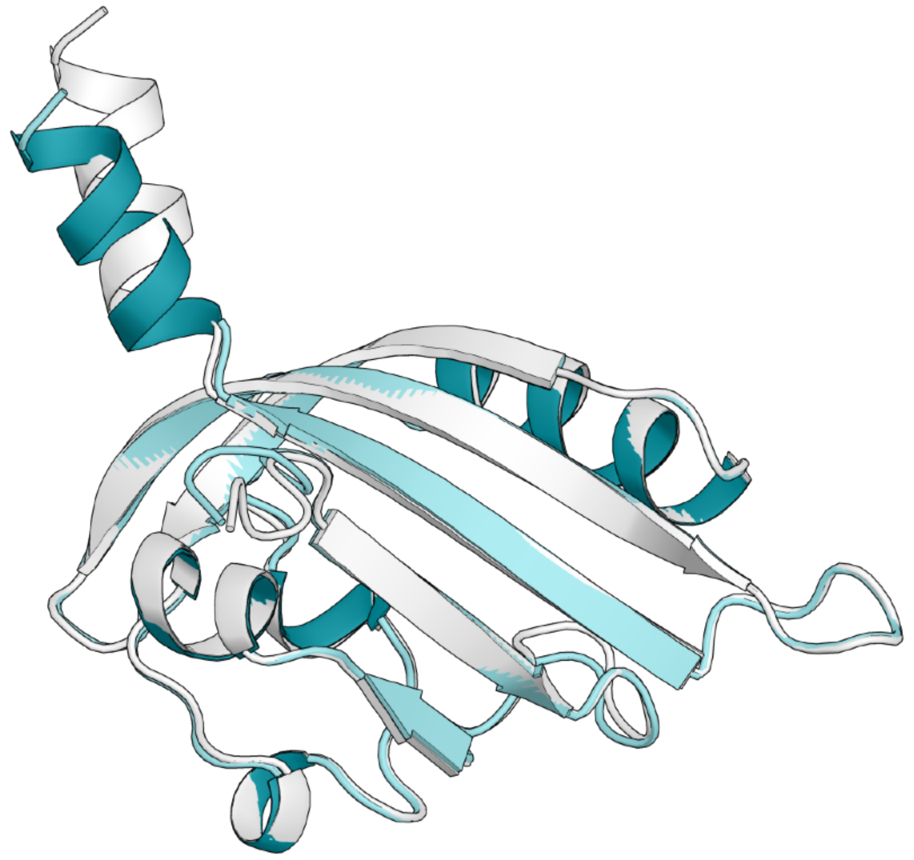

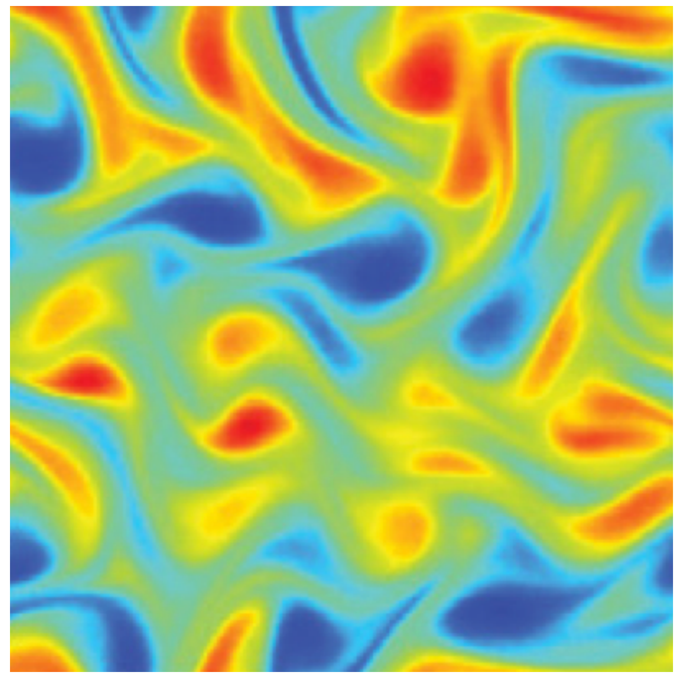

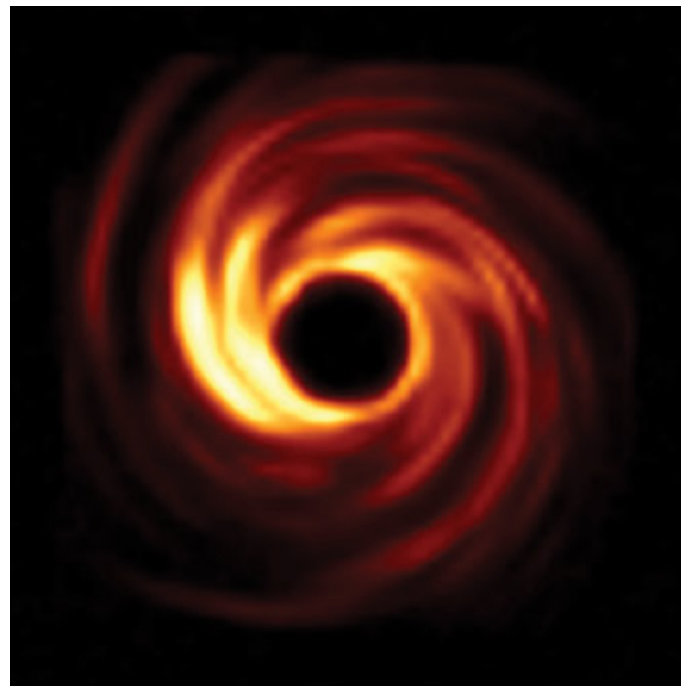

[Zheng et al., 2025]

Ω **valid design space**

DISCRETE (e.g., $\Omega = \{a_1, \ldots, a_n\}$)

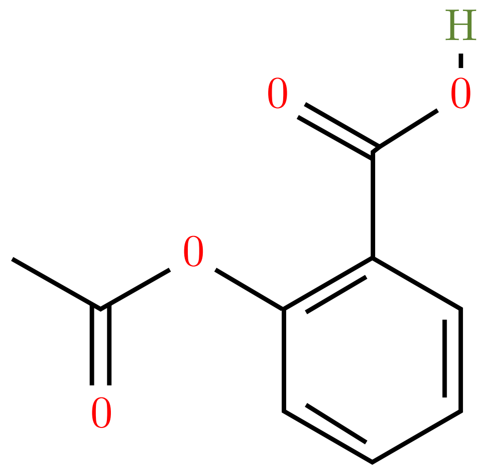

CONTINUOUS (e.g., $\Omega = [0, 1]^d$)

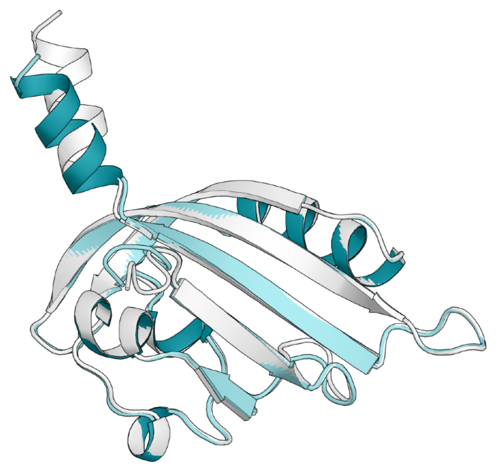

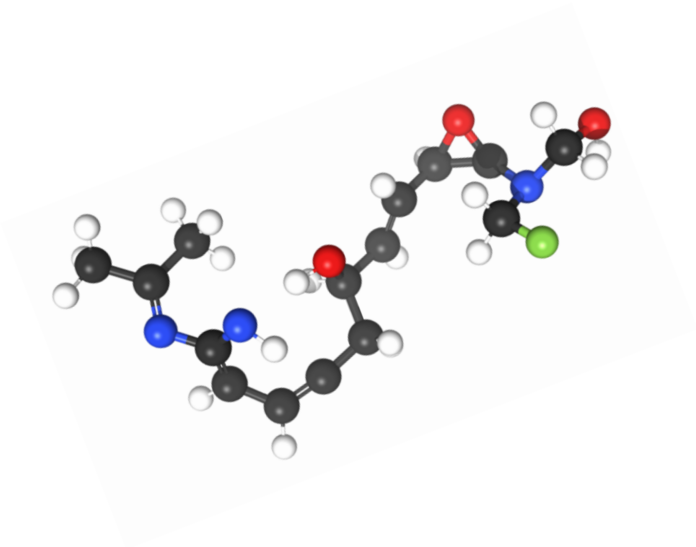

Ω **valid design space**

DISCRETE (e.g., $\Omega = \{a_1, \ldots, a_n\}$)

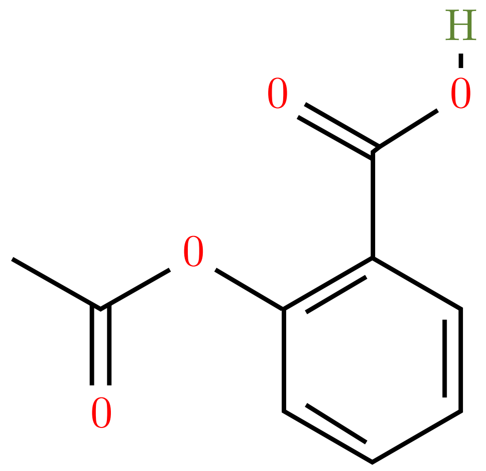

CONTINUOUS (e.g., $\Omega = [0,1]^d$)

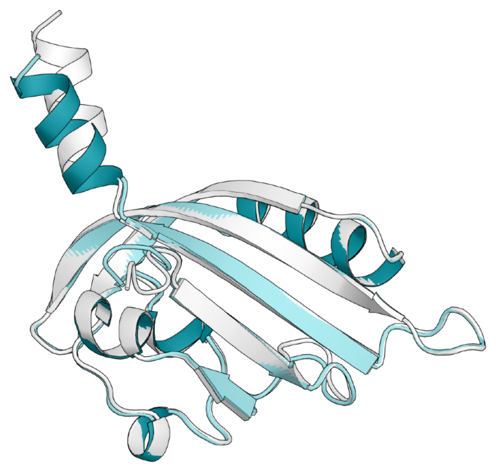

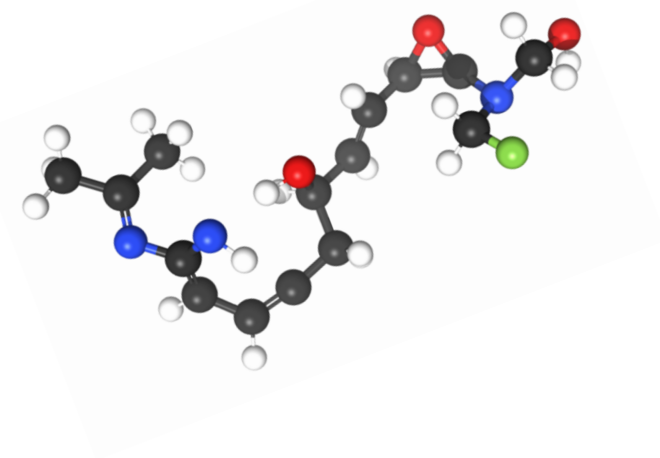

$\approx 10^{60}$ drug discovery

$\approx 10^{80}$ material design

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

Ω **valid design space**

DISCRETE (e.g., $\Omega = \{a_1, \ldots, a_n\}$)

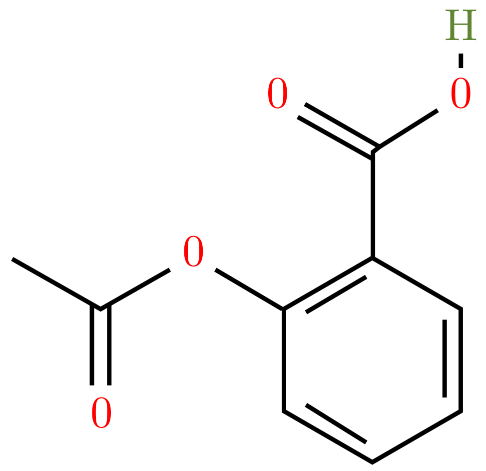

CONTINUOUS (e.g., $\Omega = [0, 1]^d$)

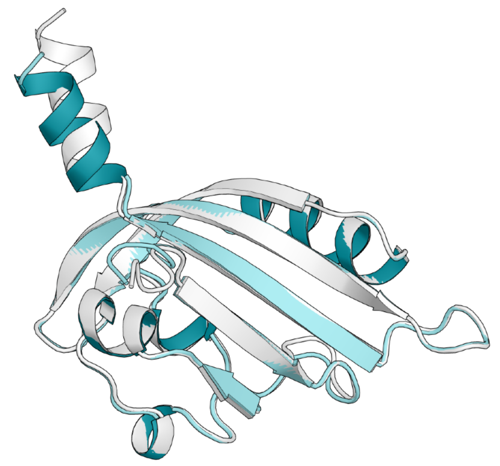

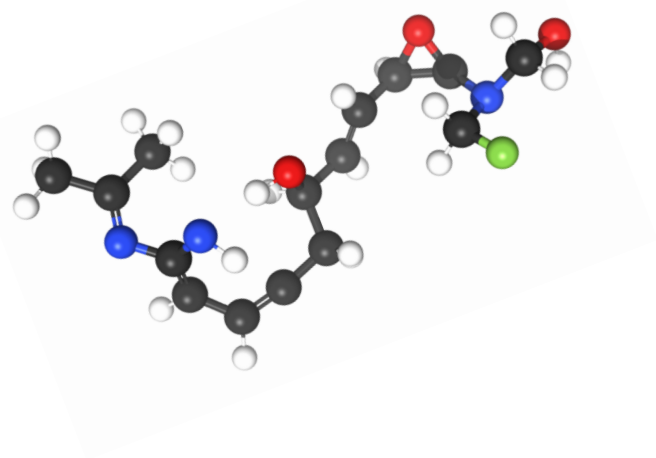

$\approx 10^{60}$ drug discovery

$\approx 10^{80}$ material design

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

→ Generative Discovery

Generative Discovery Paradigm

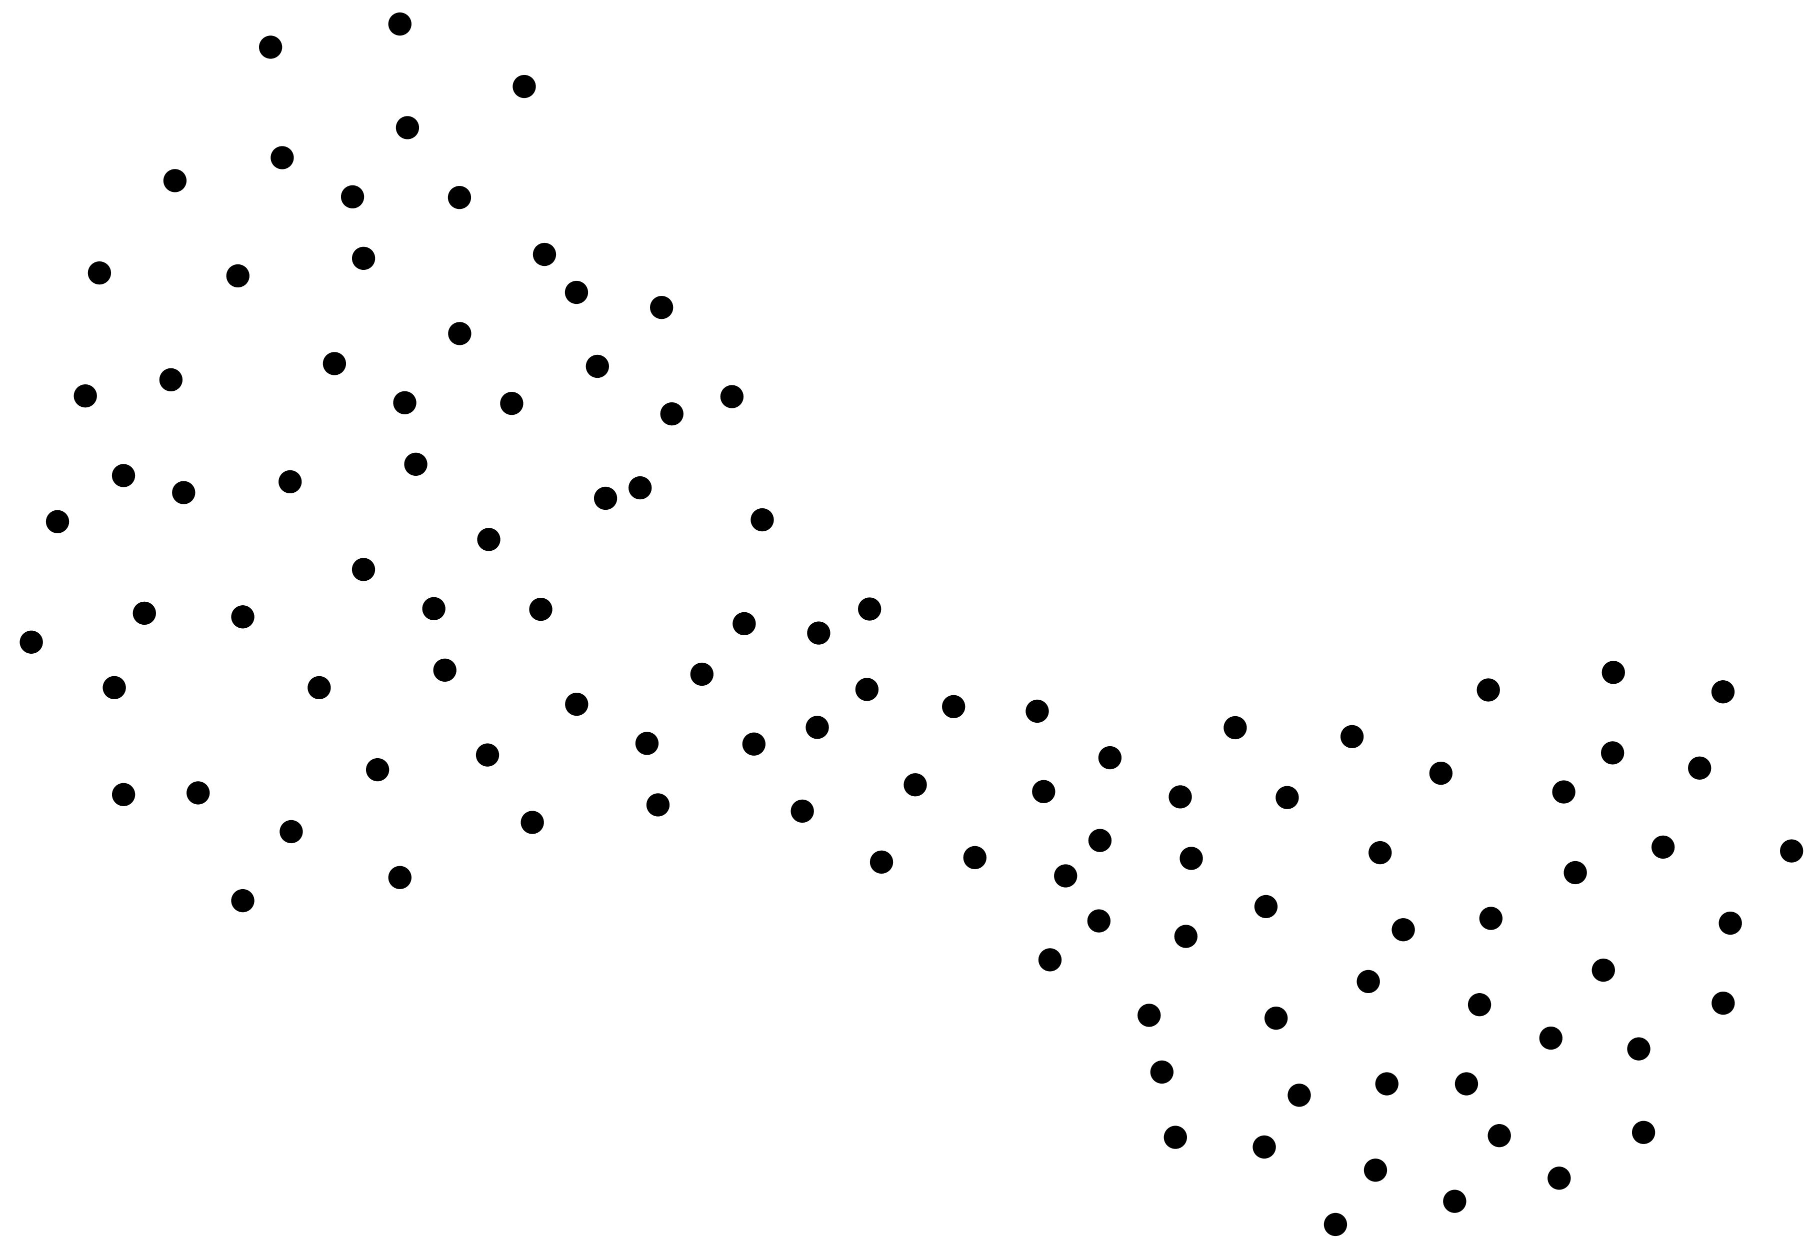

Data $D = \{x_i\}_{i=1}^n$

Generative Discovery Paradigm

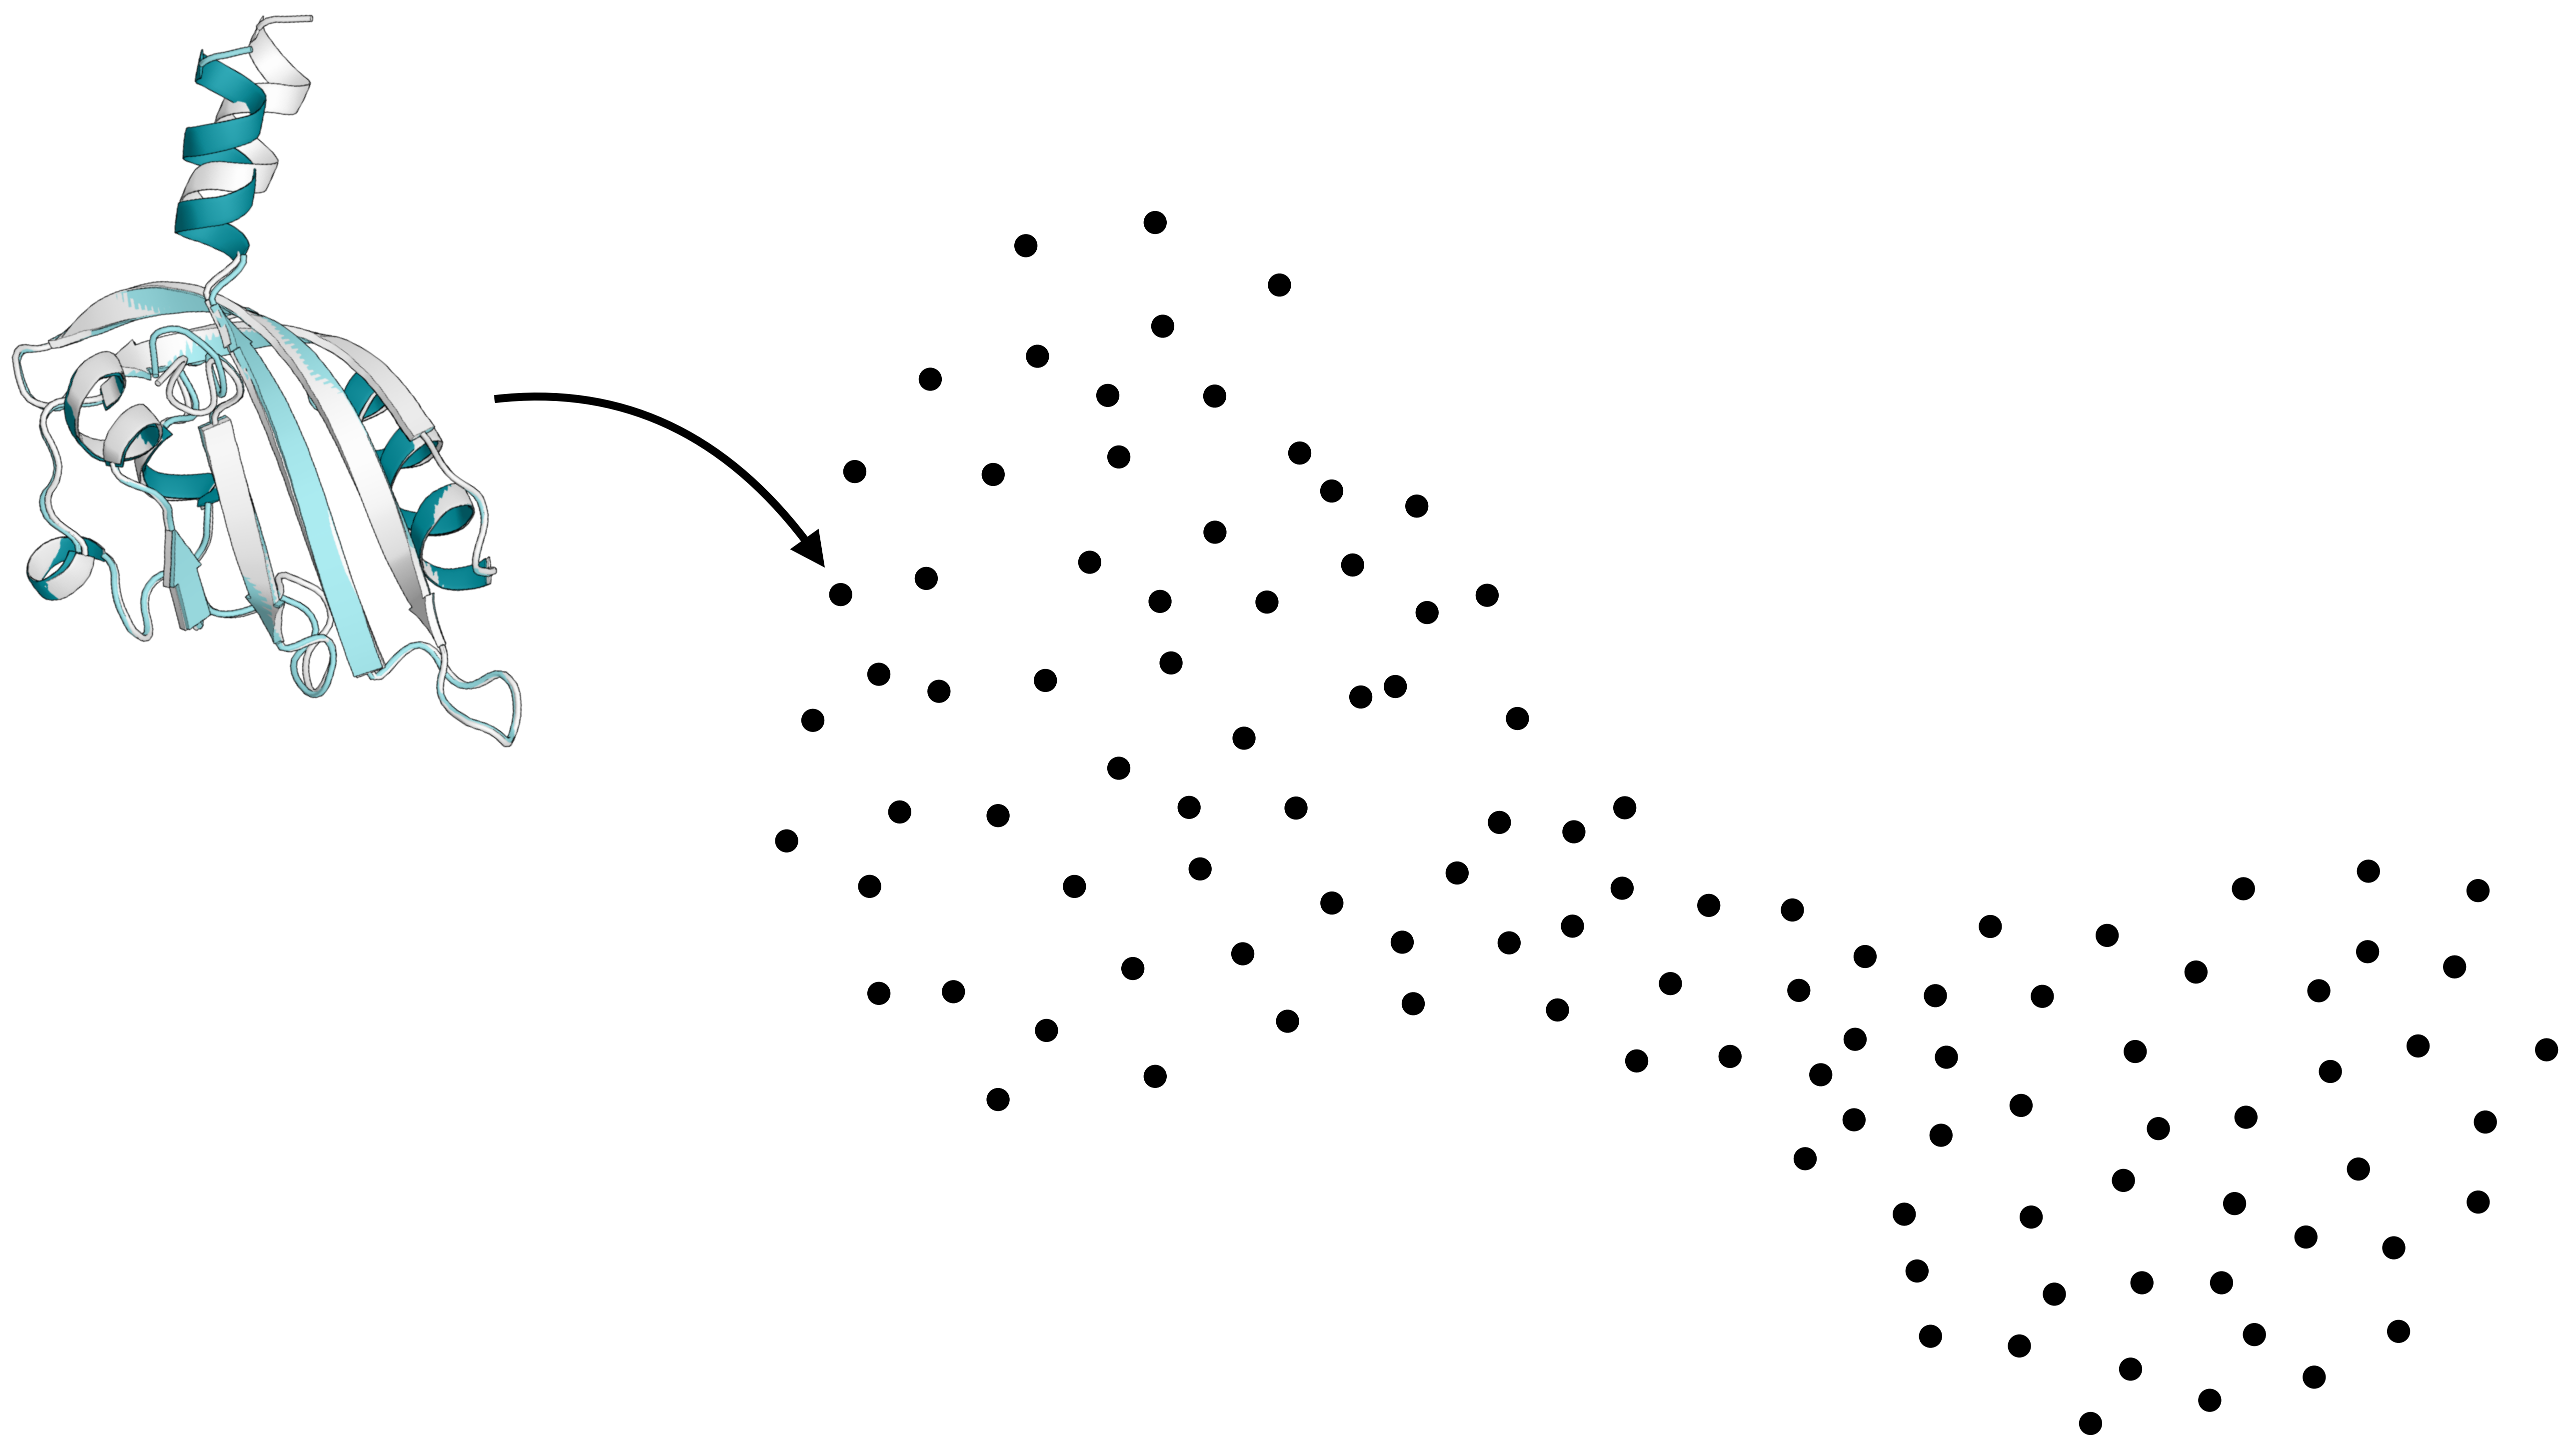

Data $D = \{x_i\}_{i=1}^n$

Generative Discovery Paradigm

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

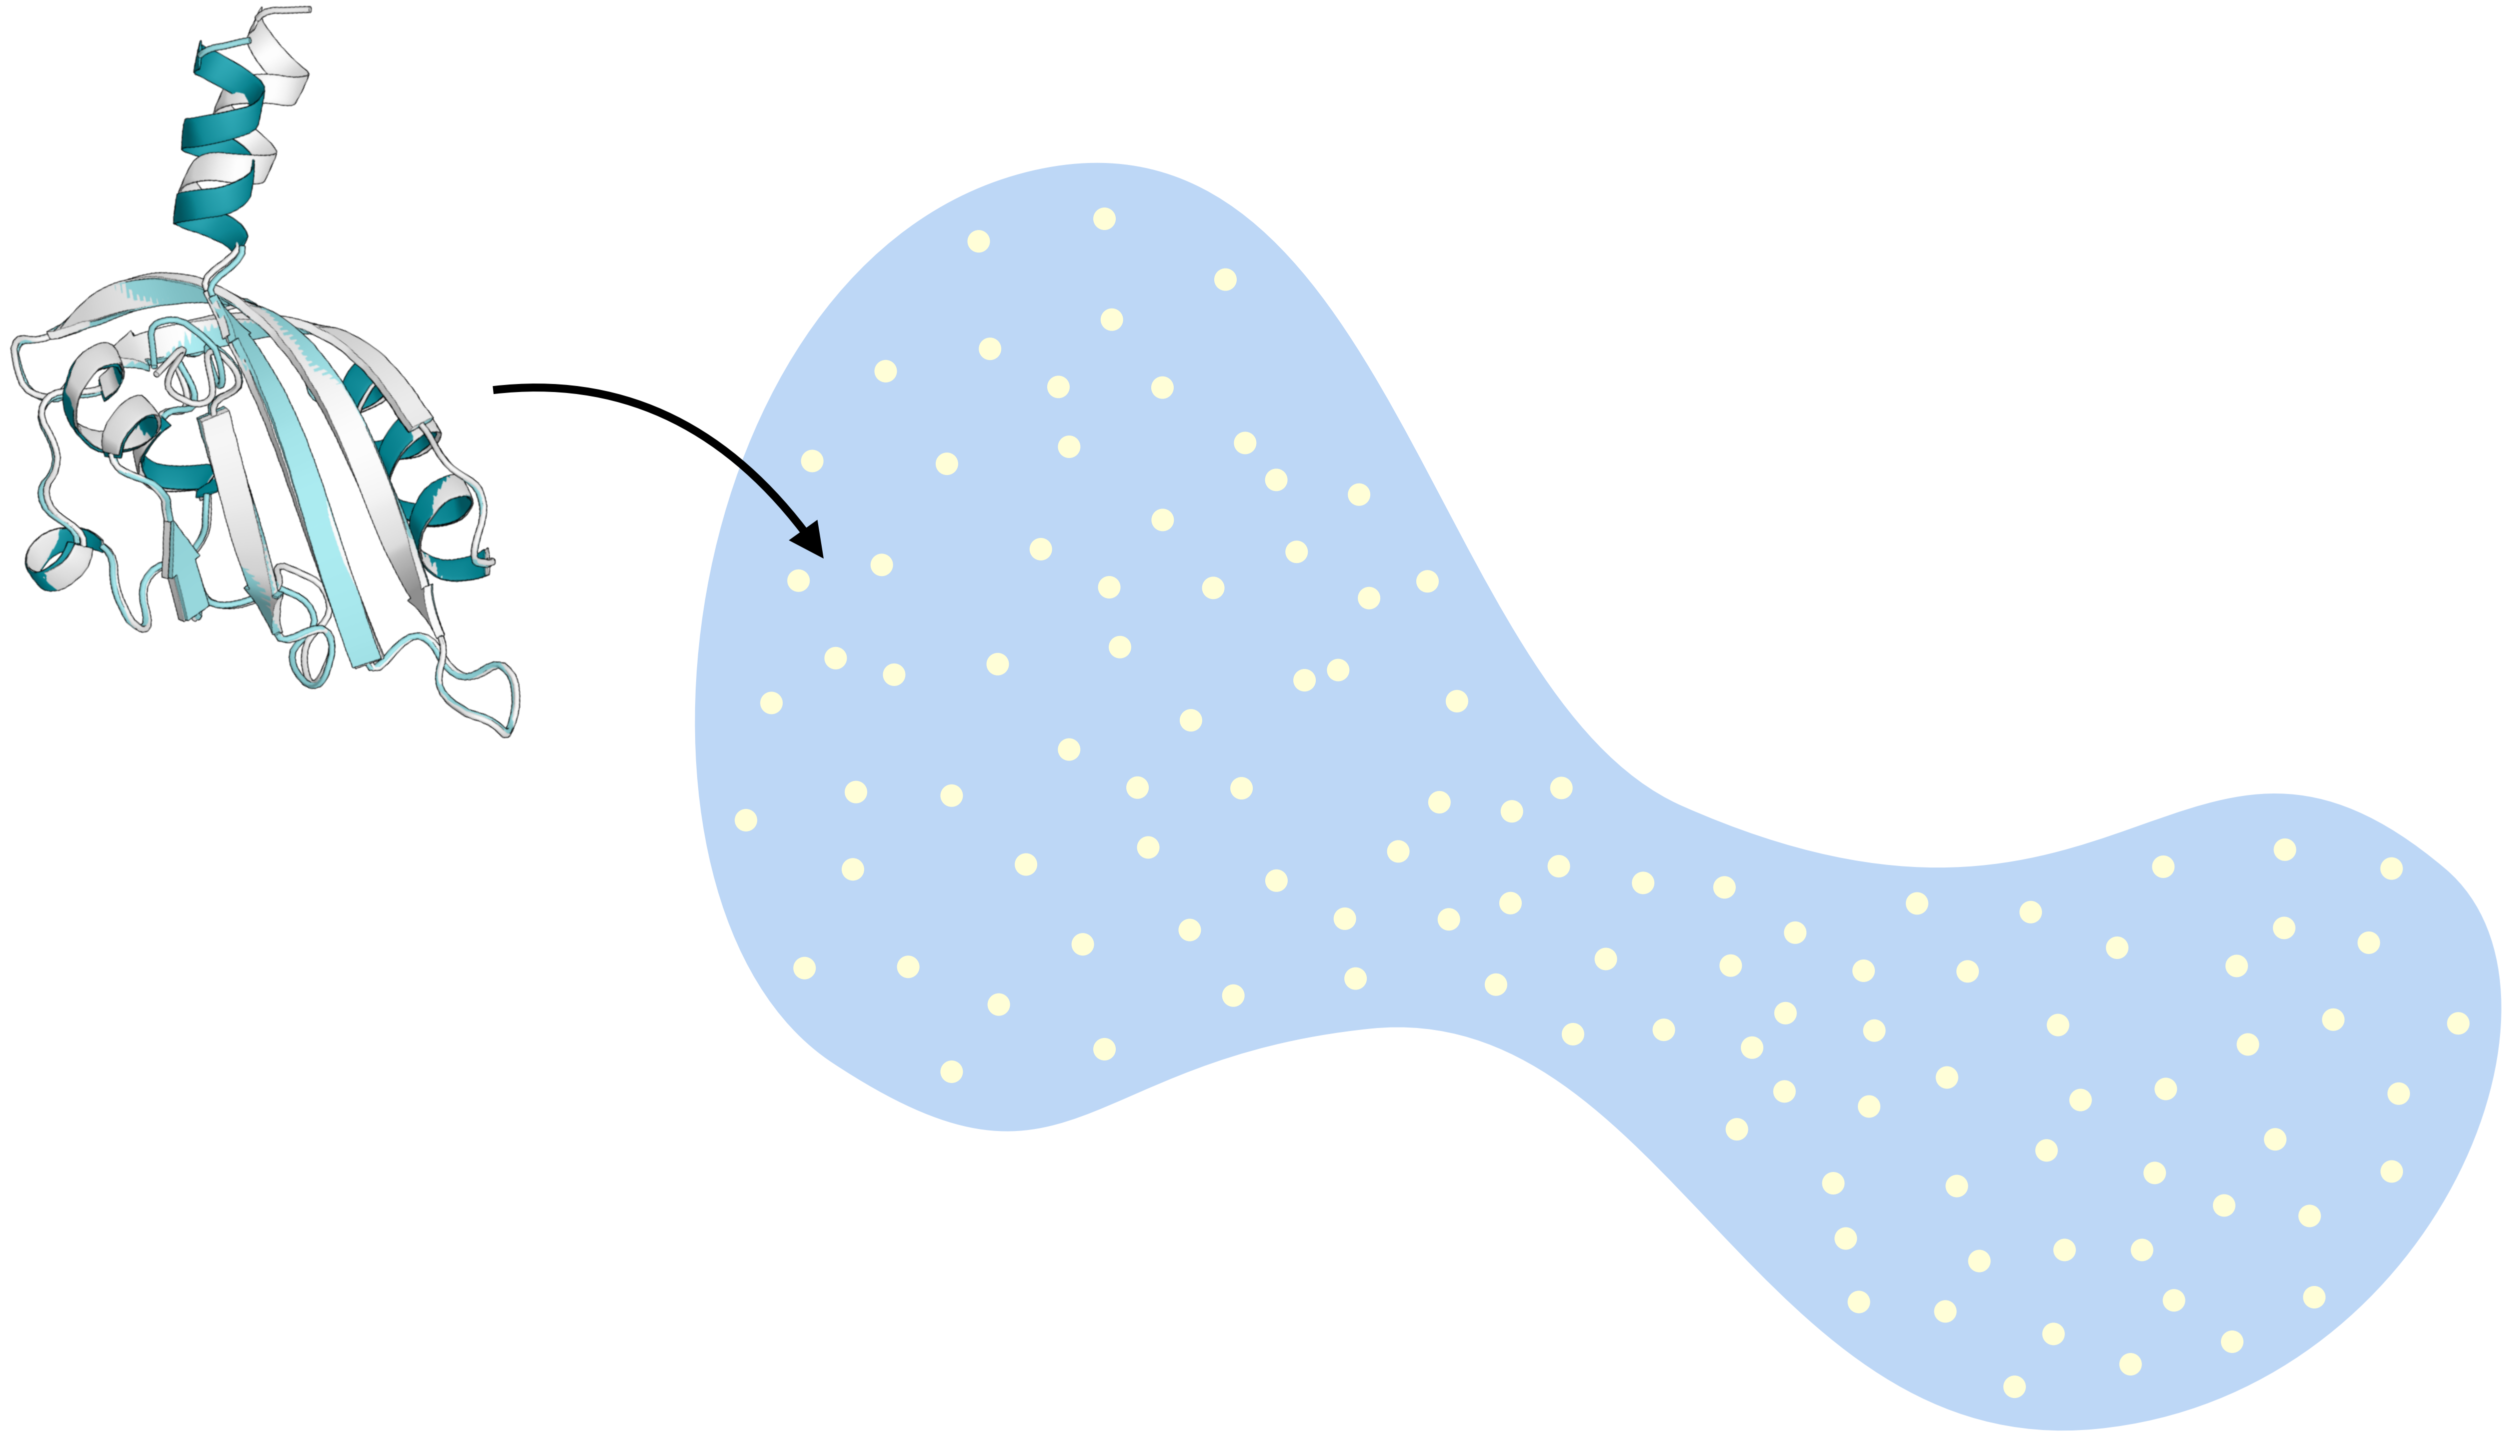

Data $D = \{x_i\}_{i=1}^n$

Generative Discovery Paradigm

valid design space Ω generable set of pre-trained gen. model Ω_{pre}

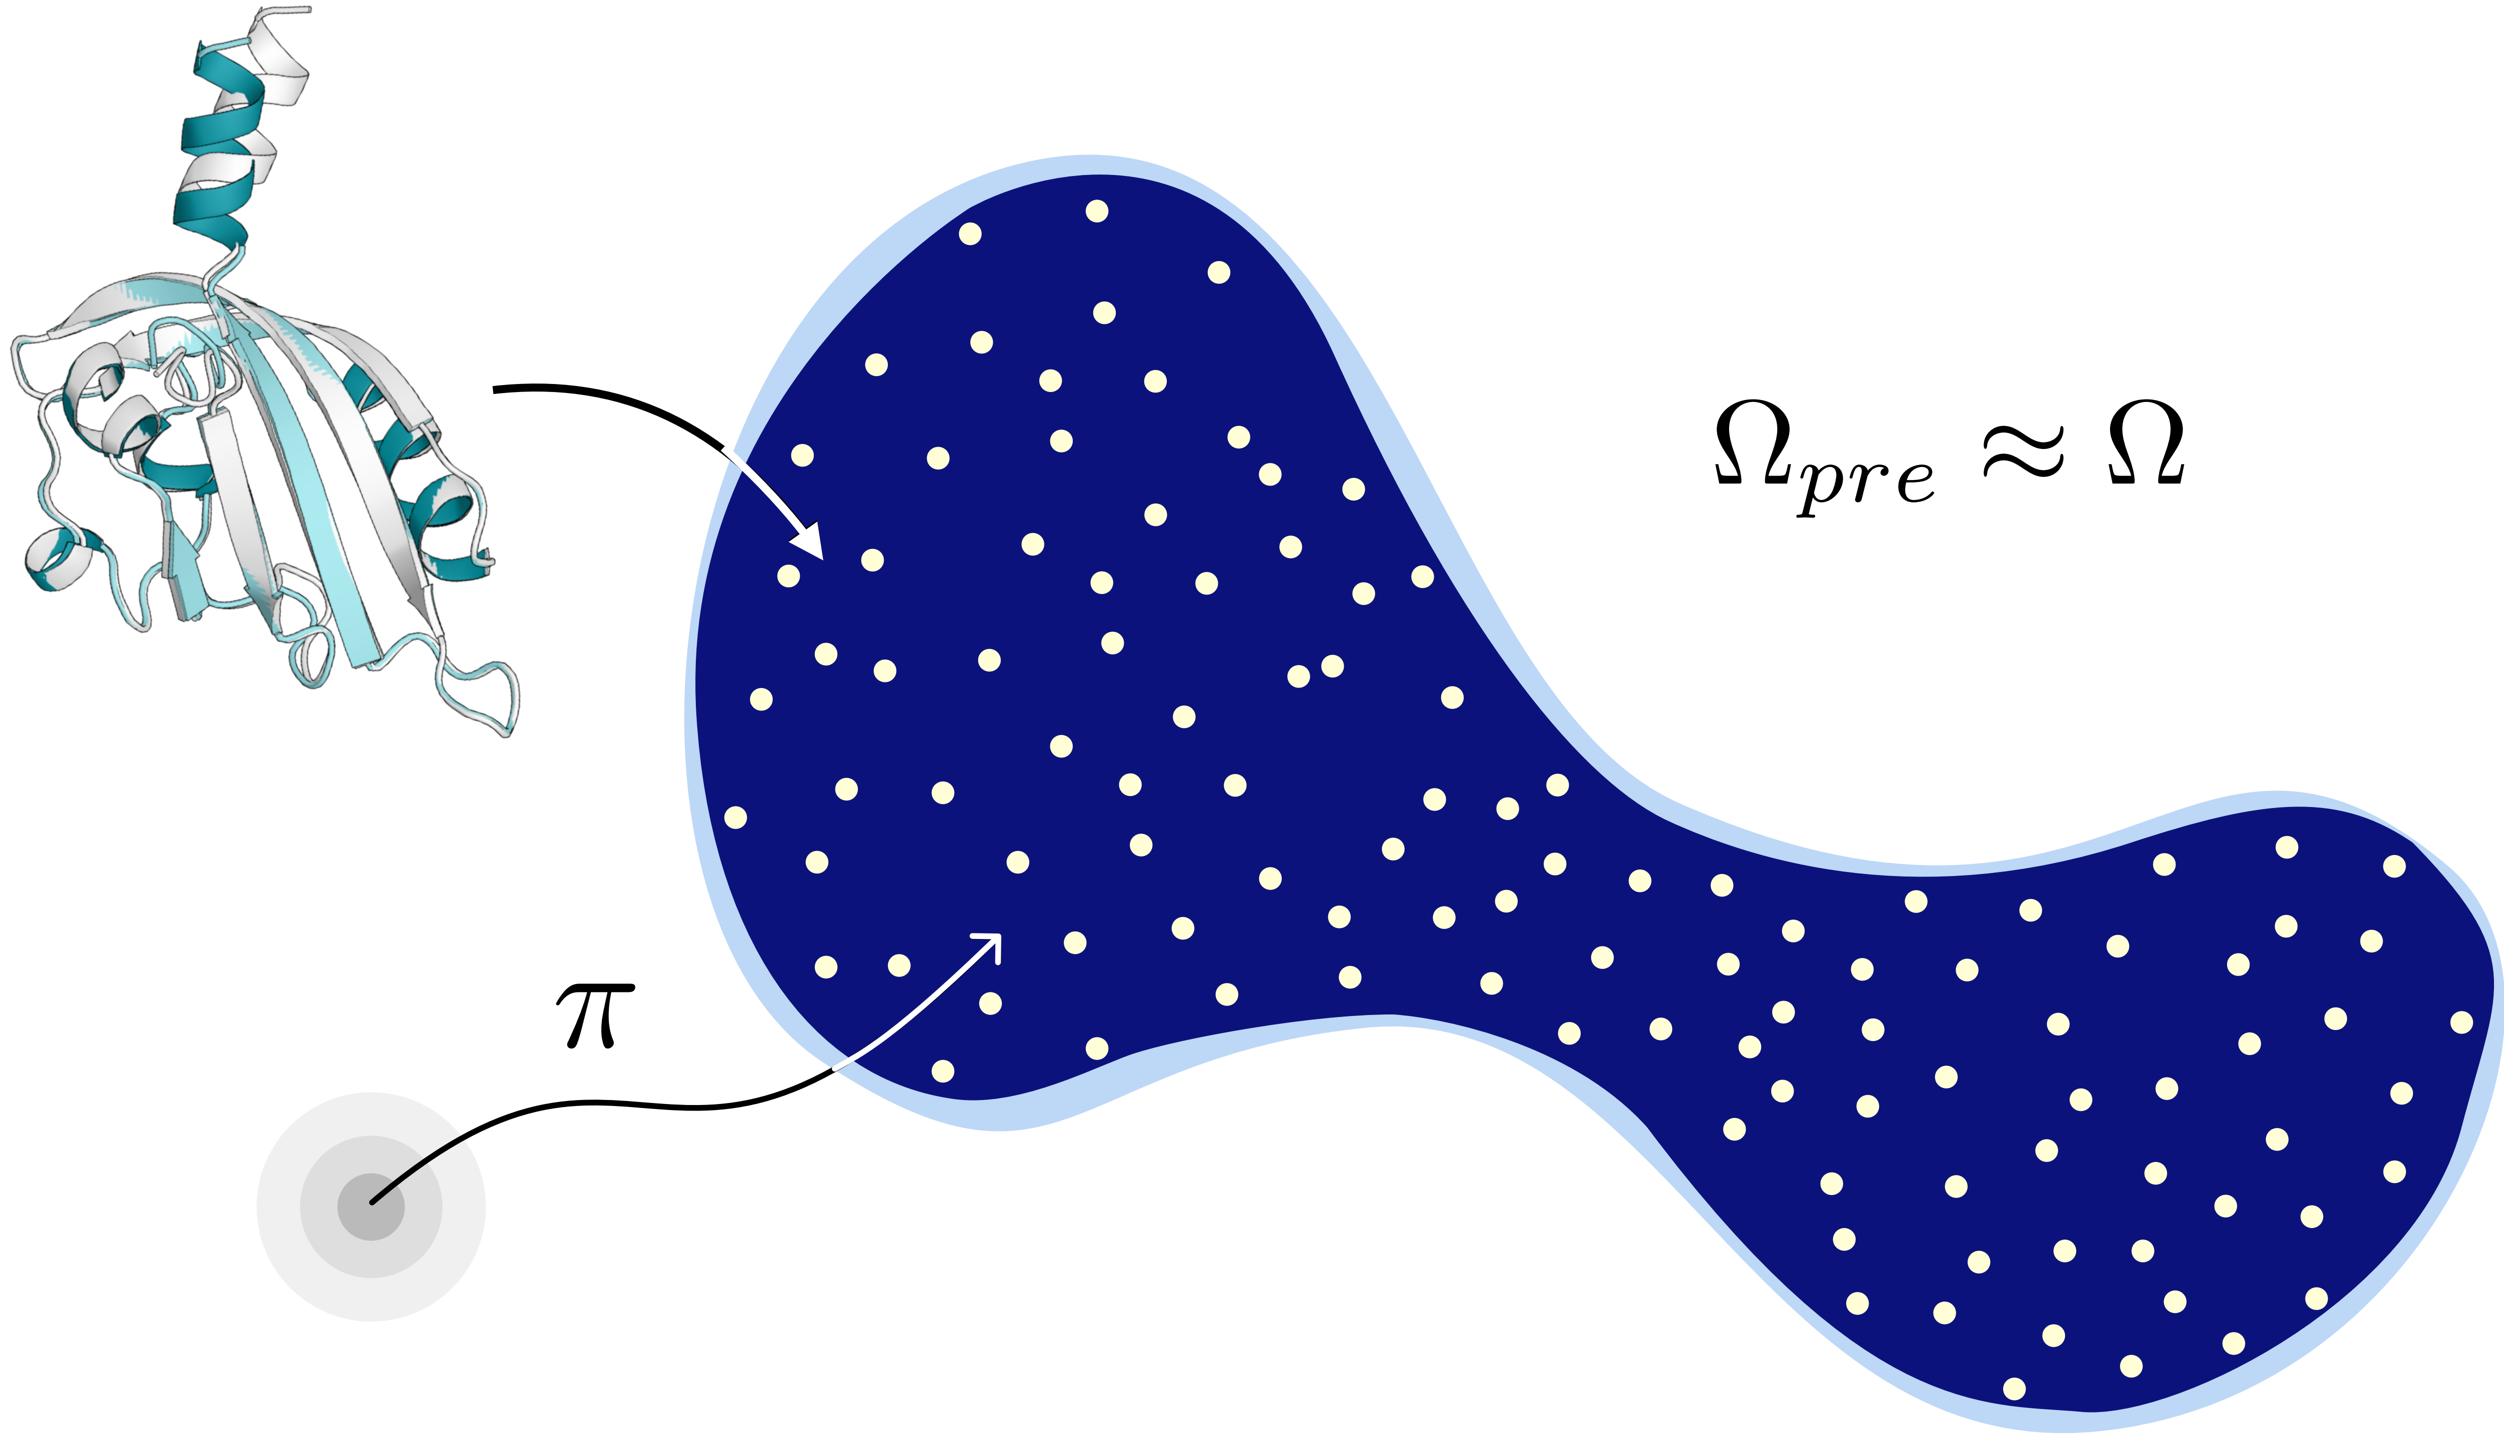

Data $D = \{x_i\}_{i=1}^n$

↓

Pre-trained generative model π

The Dream: Discovery via Generative Optimization

valid design space Ω

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

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

↓

Pre-trained generative model π

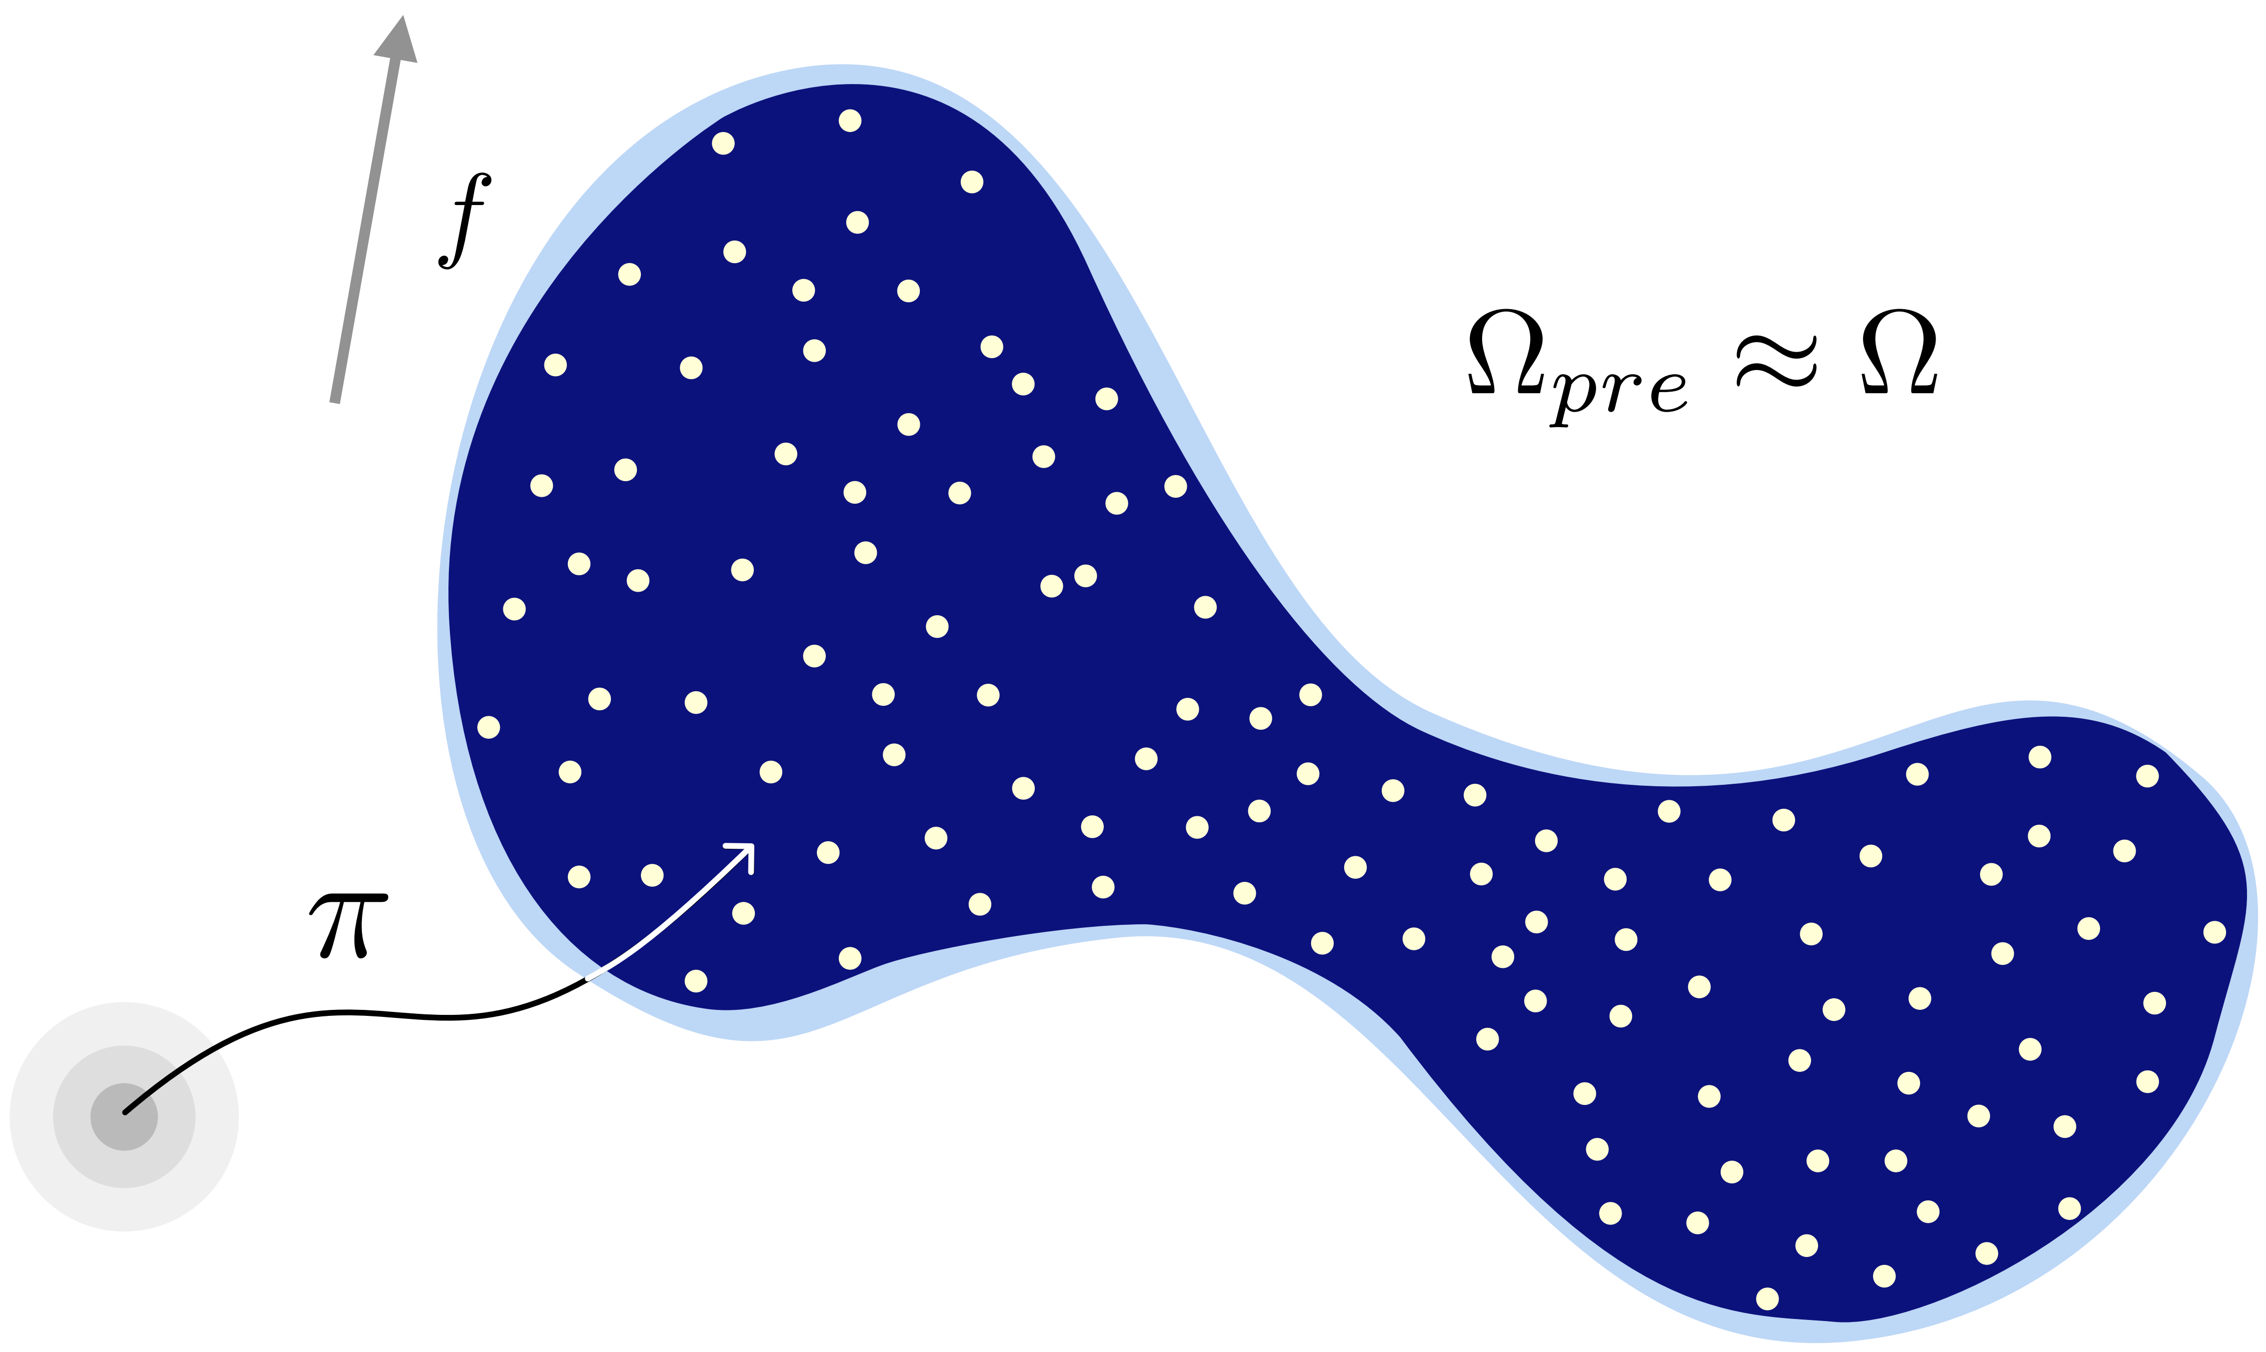

Discovery problem

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

The Dream: Discovery via Generative Optimization

valid design space Ω — generable set of pre-trained gen. model Ω_{pre}

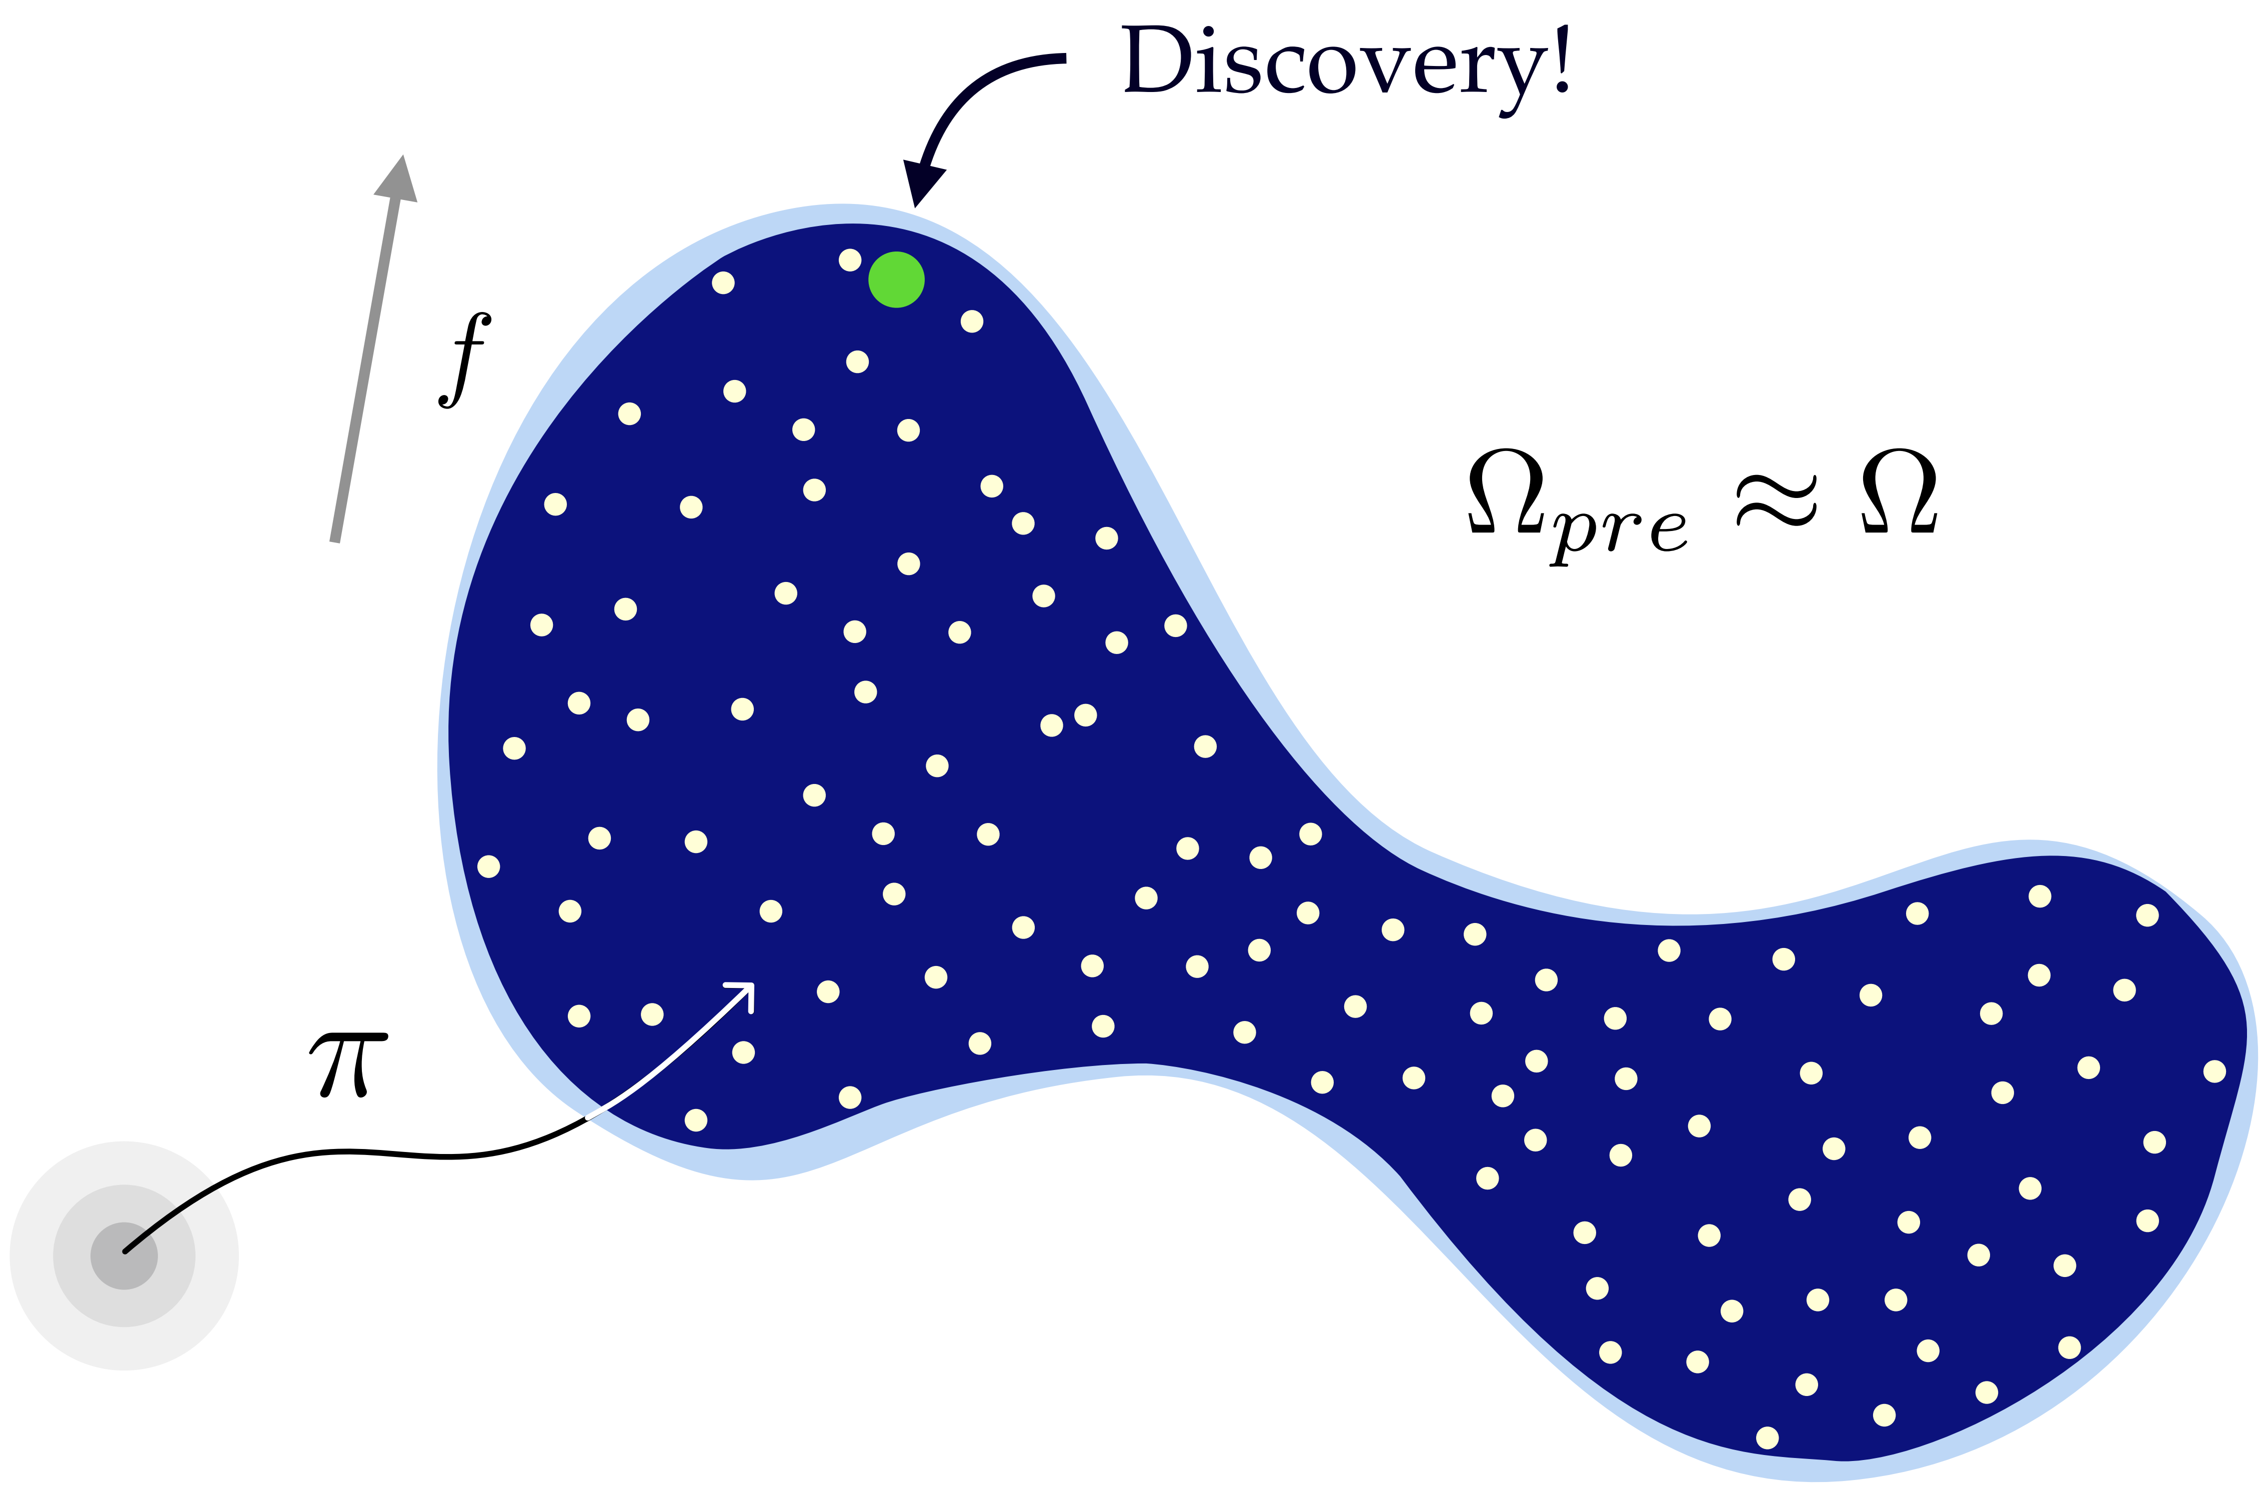

Data $D = \{x_i\}_{i=1}^n$

↓

Pre-trained generative model π

↓

Adapt π to optimize f

Discovery problem

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

From Generative Modeling to Generative Discovery

Generative Modeling

Learn generative prior from data

From Generative Modeling to Generative Discovery

Generative Modeling

Learn generative prior from data

↓

Current Status

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

From Generative Modeling to Generative Discovery

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 events (structures)** of **exceptionally high-value** and **far beyond existing data**

This talk:

Foundations of Generative Discovery Beyond the Data

Part I

Tail-Aware Reward Adaptation

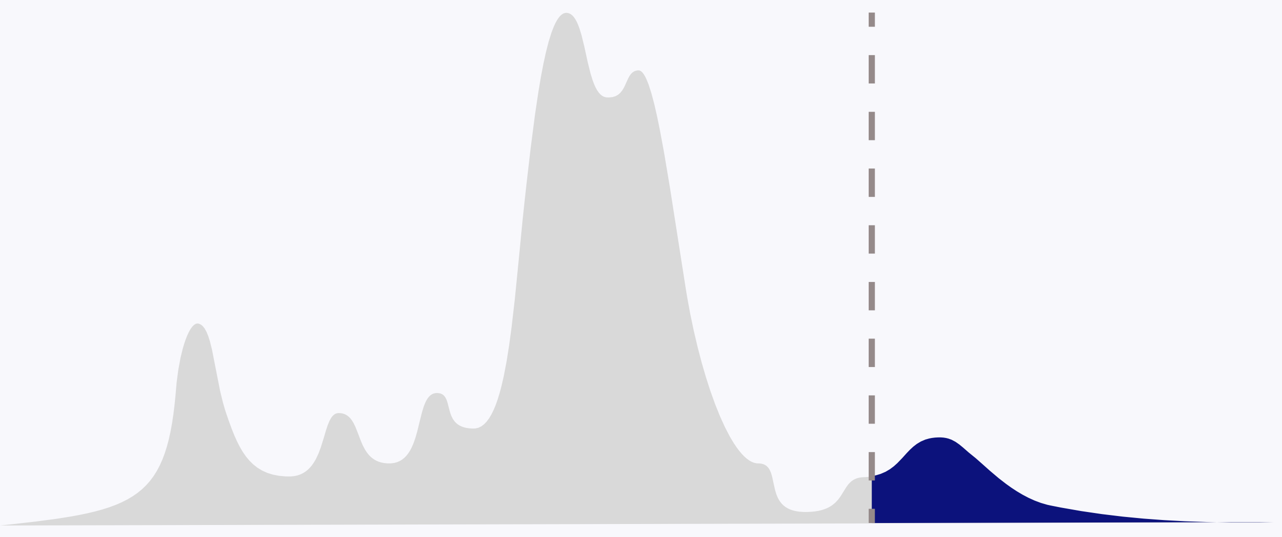

$\mathrm{CVaR}^{f}_{1-\beta}(p_1^{\pi})$

Part II

Data Debiasing and Hidden Mode Discovery

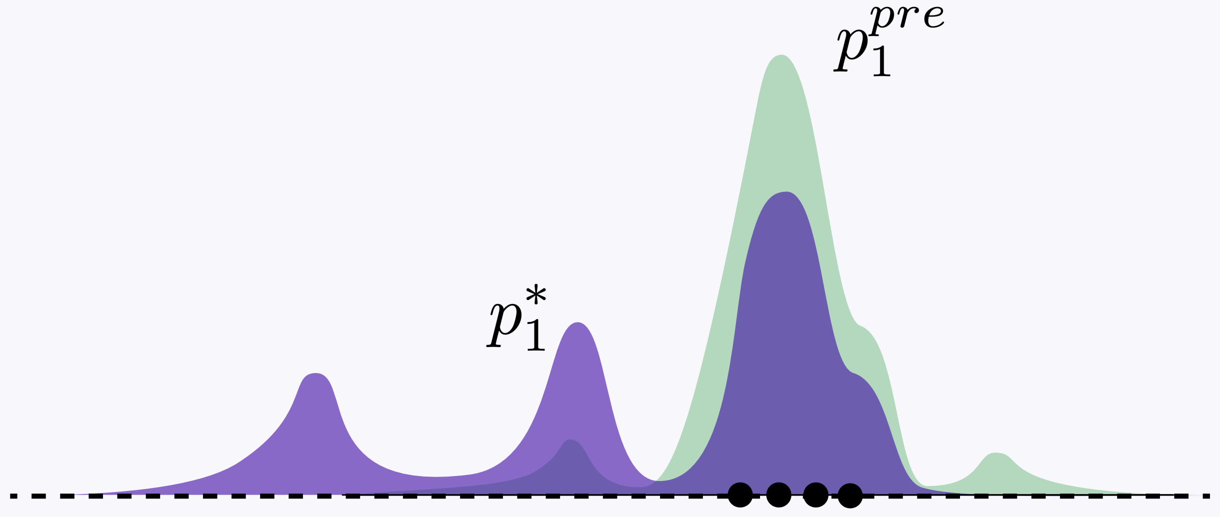

Part III

Out-of-Distribution Flow Modeling

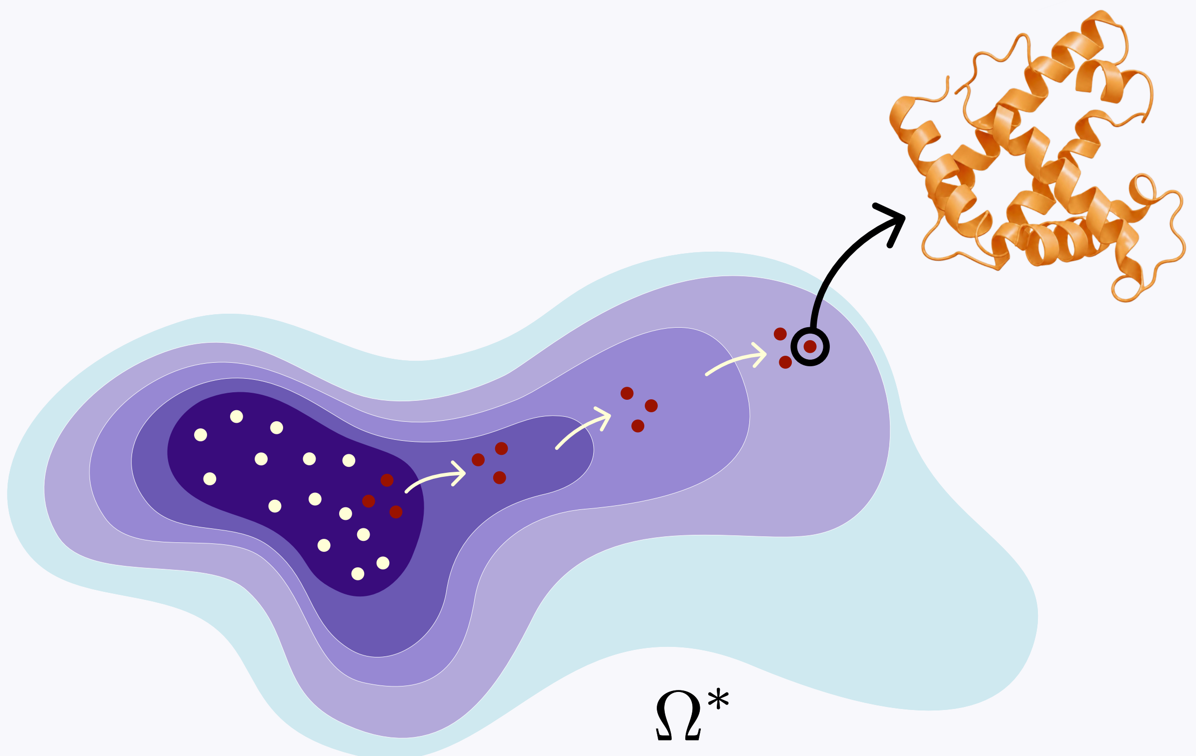

This talk:

Foundations of Generative Discovery Beyond the Data

Part I

Tail-Aware Reward Adaptation

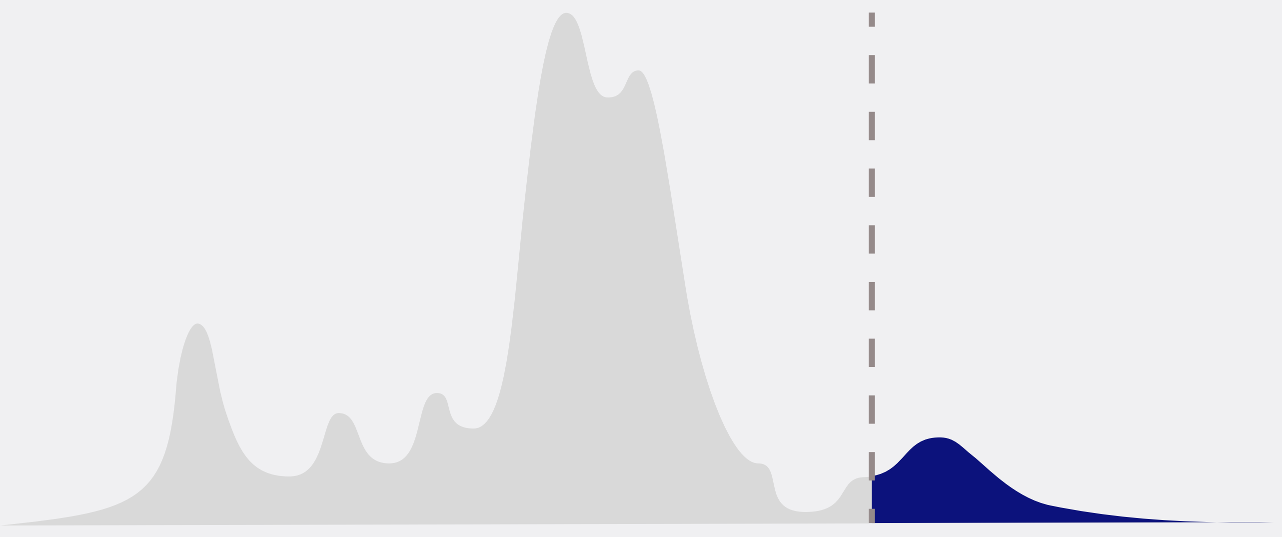

Part II

Data Debiasing and Hidden Mode Discovery

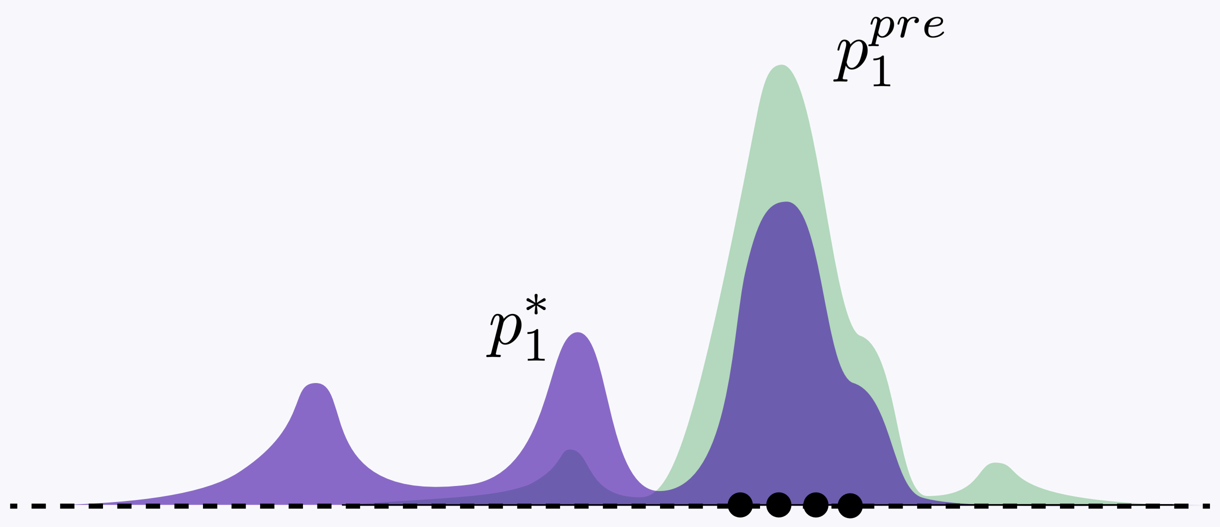

Part III

Out-of-Distribution Flow Modeling

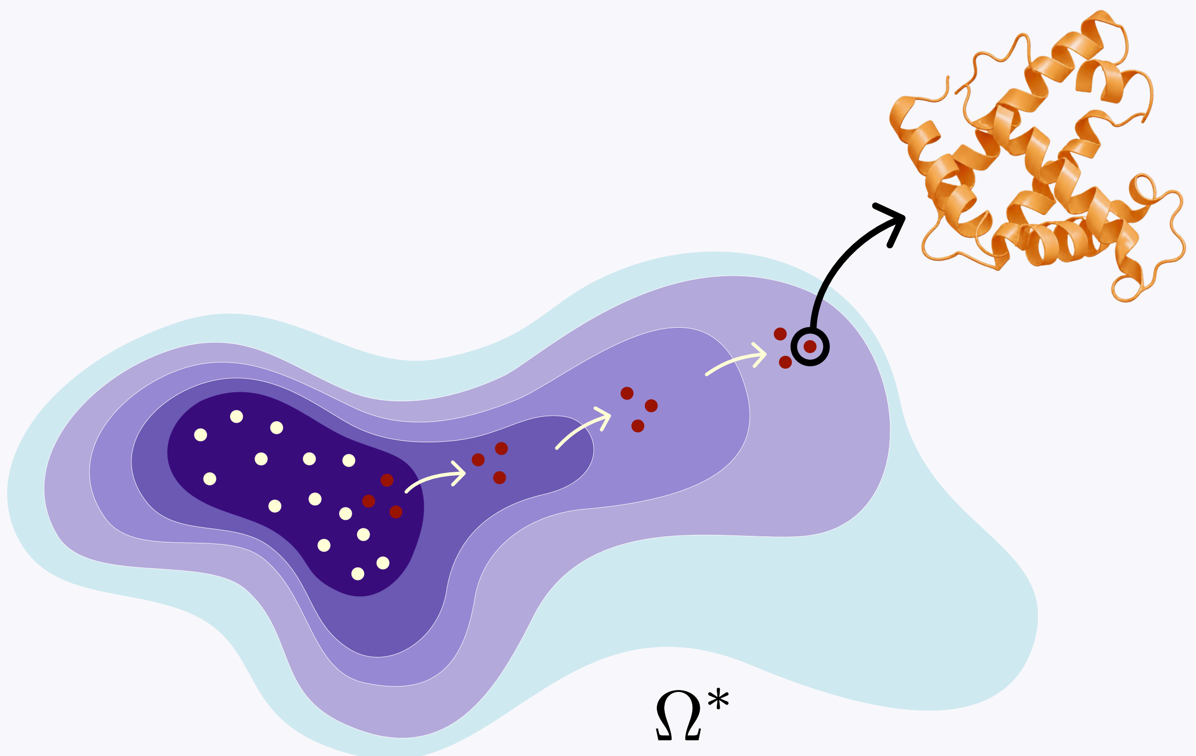

Diffusion and Flow Modeling

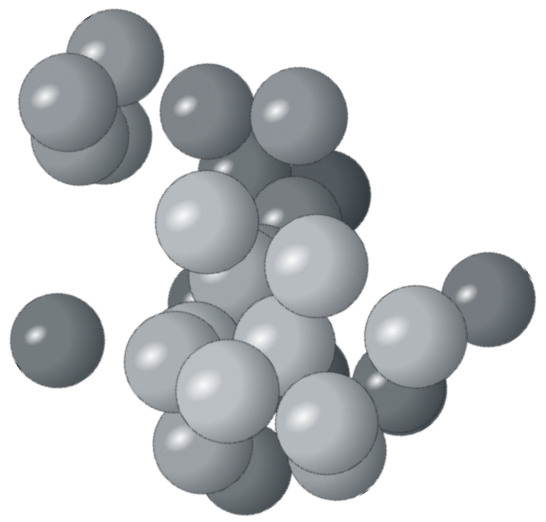

noise
distribution

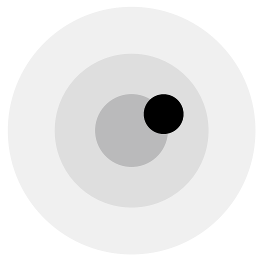

$X_0 \sim p_0$

Generation via learned vector field

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

Diffusion and Flow Modeling

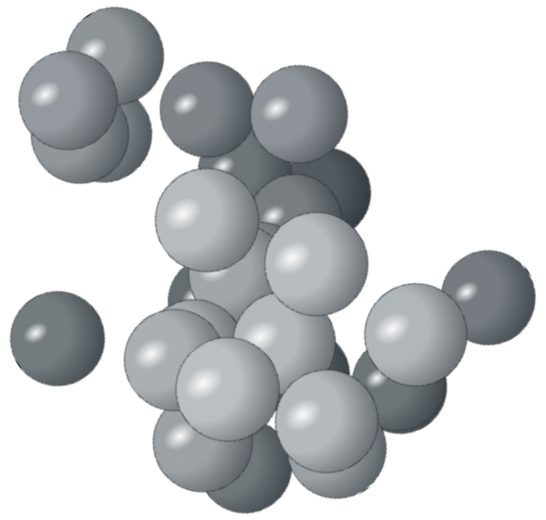

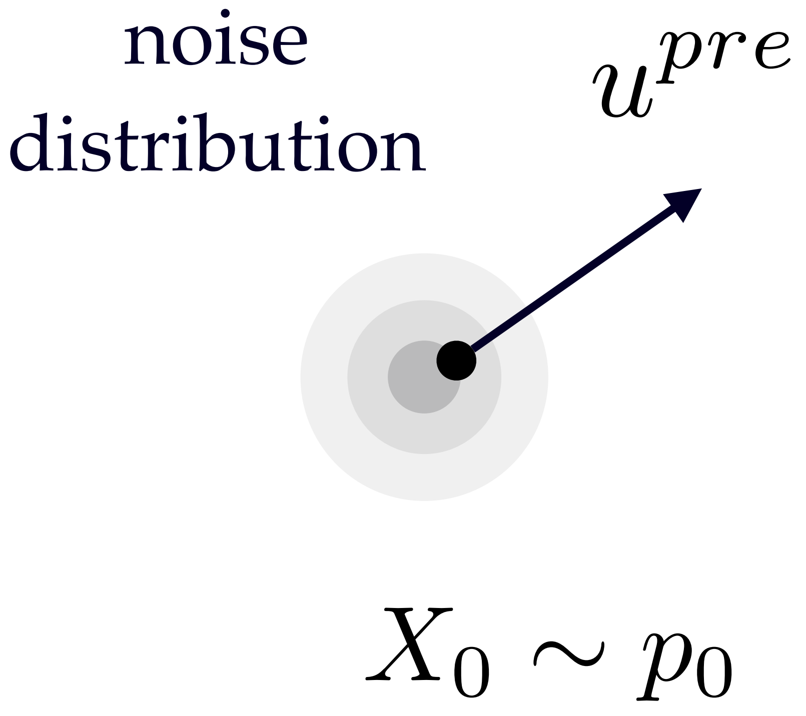

Generation via learned vector field

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

Diffusion and Flow Modeling

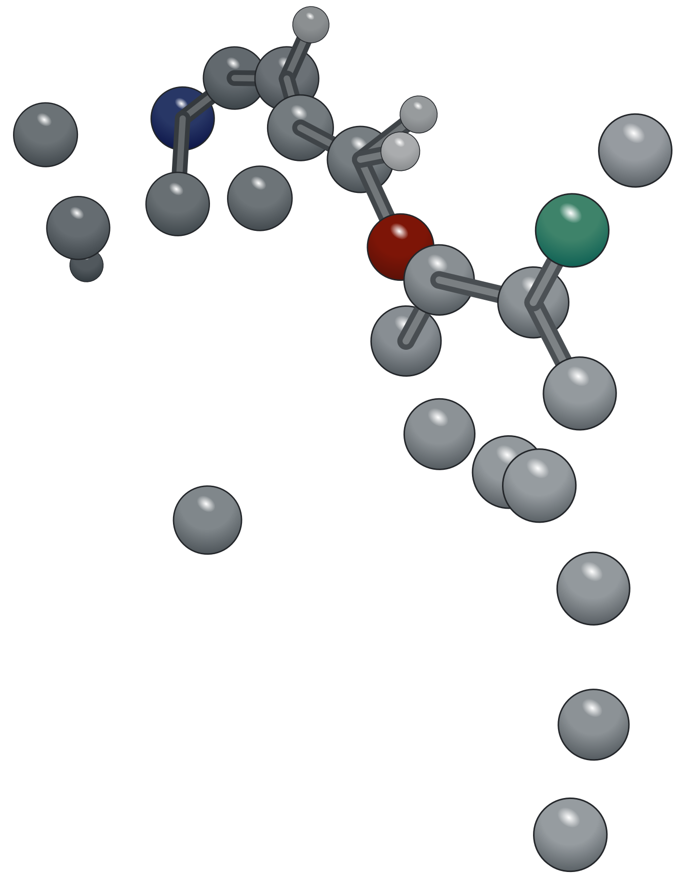

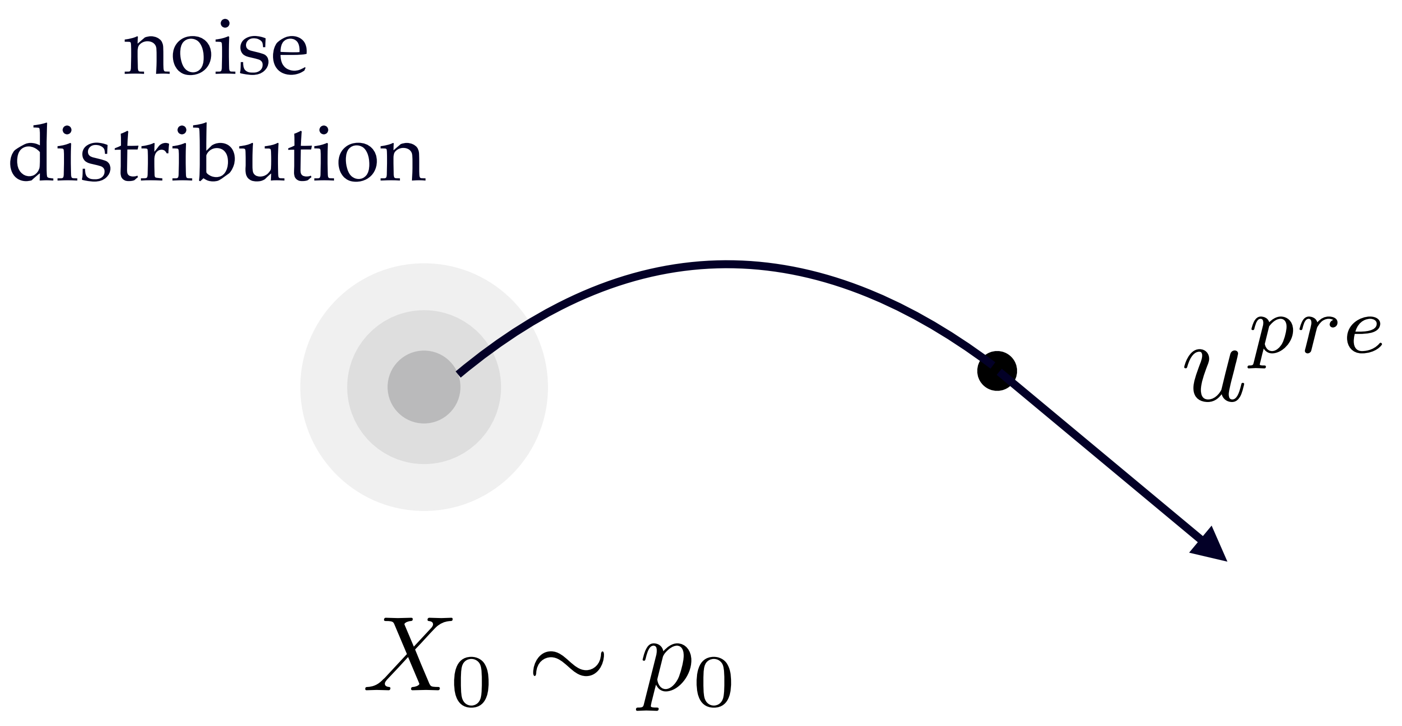

Generation via learned vector field

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

Diffusion and Flow Modeling

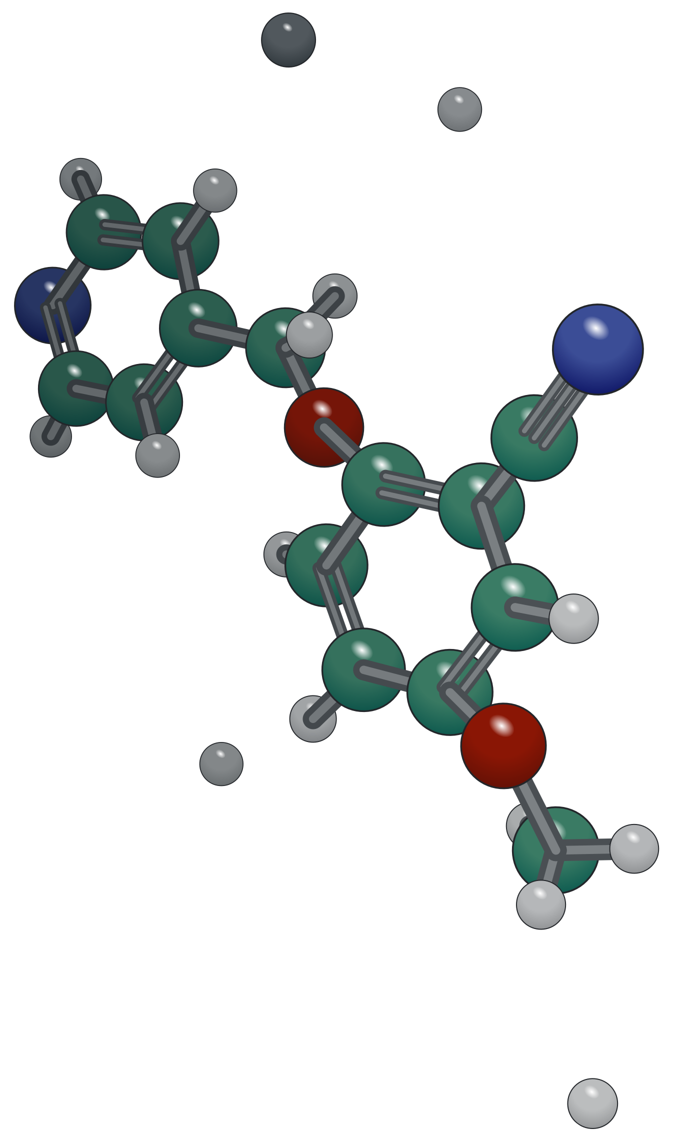

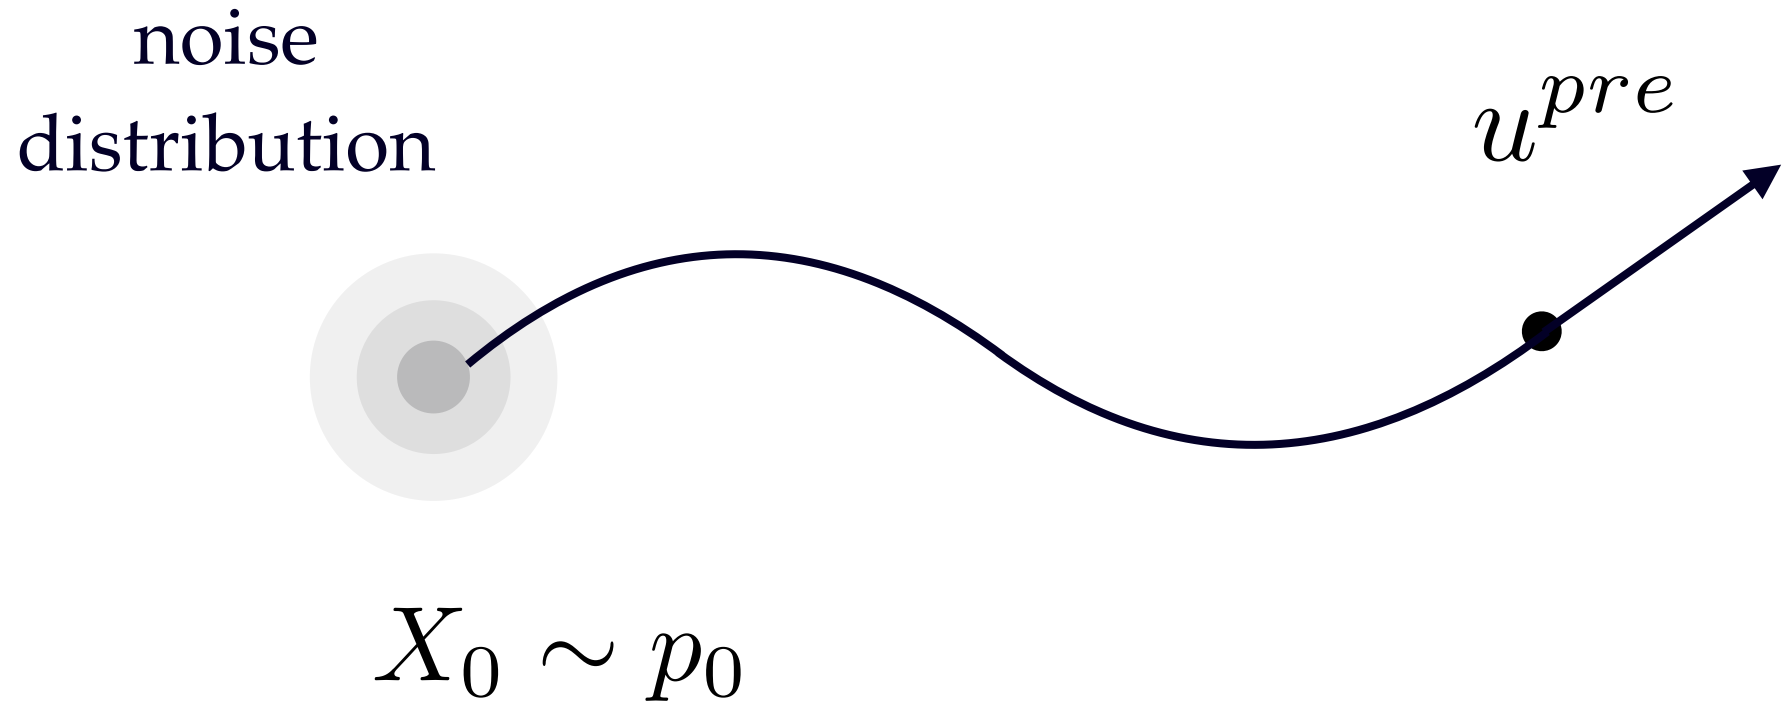

Generation via learned vector field

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

Diffusion and Flow Modeling

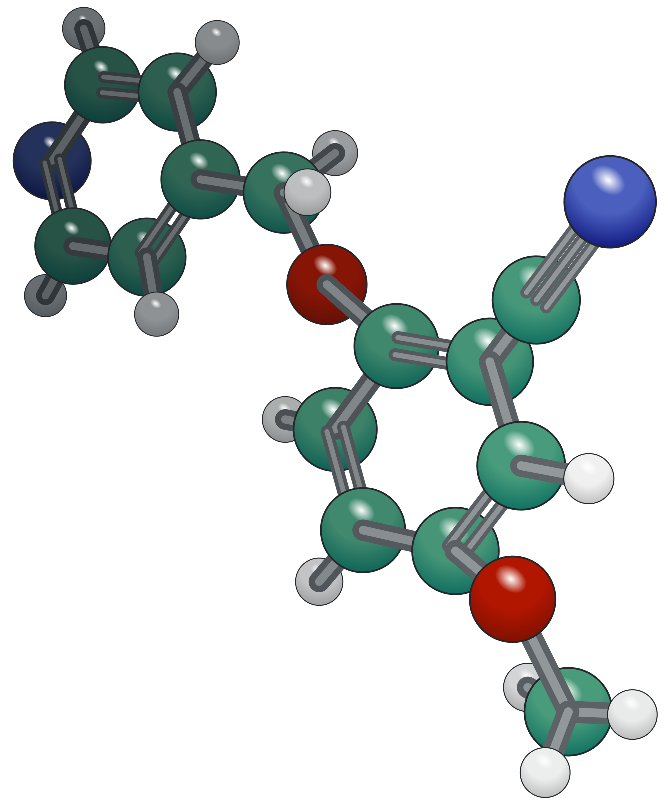

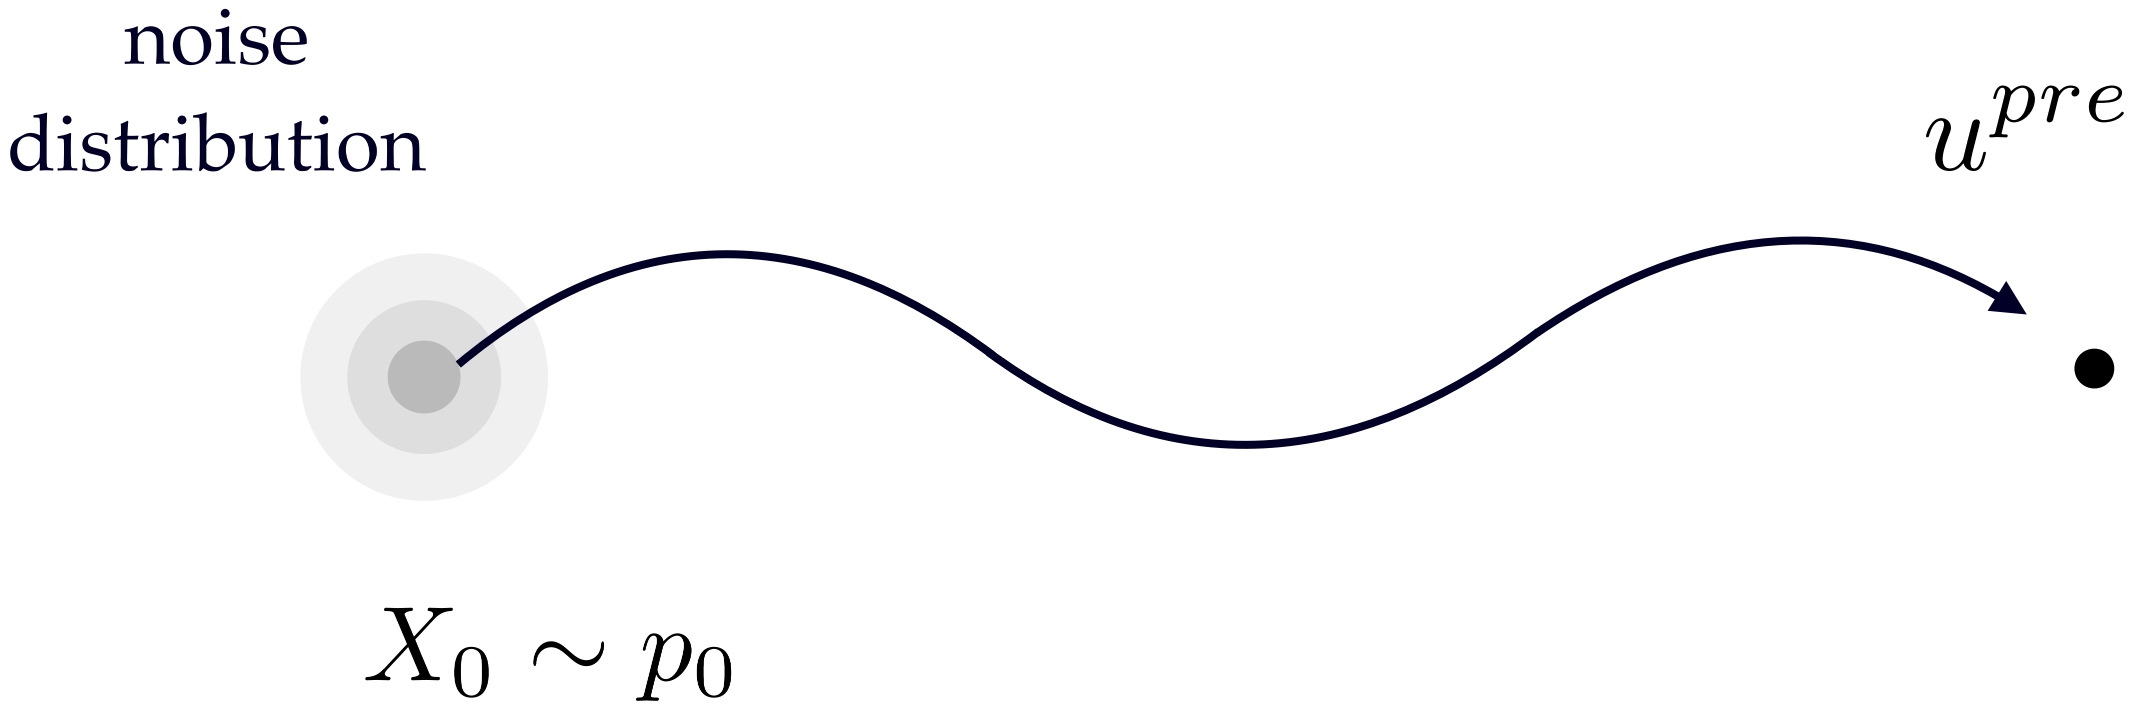

Generation via learned vector field

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

Diffusion and Flow Modeling

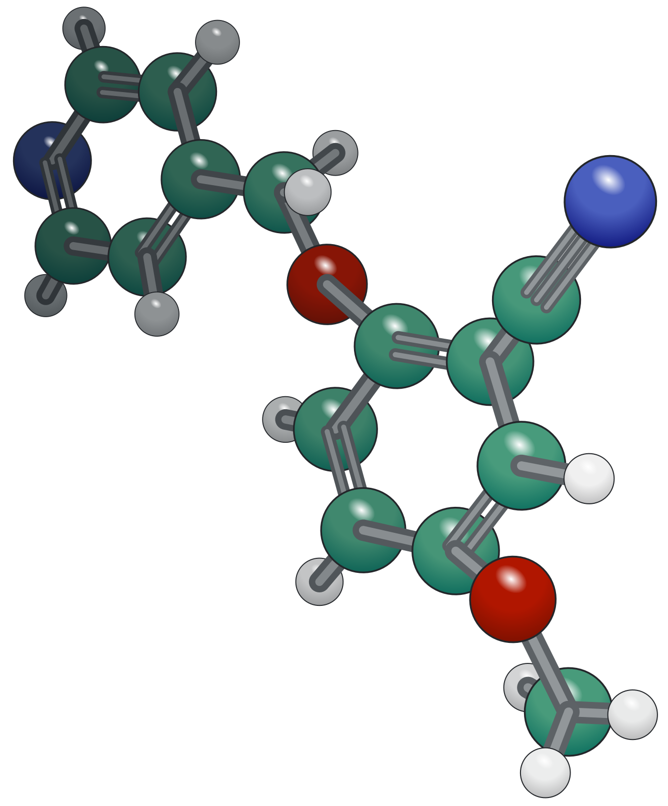

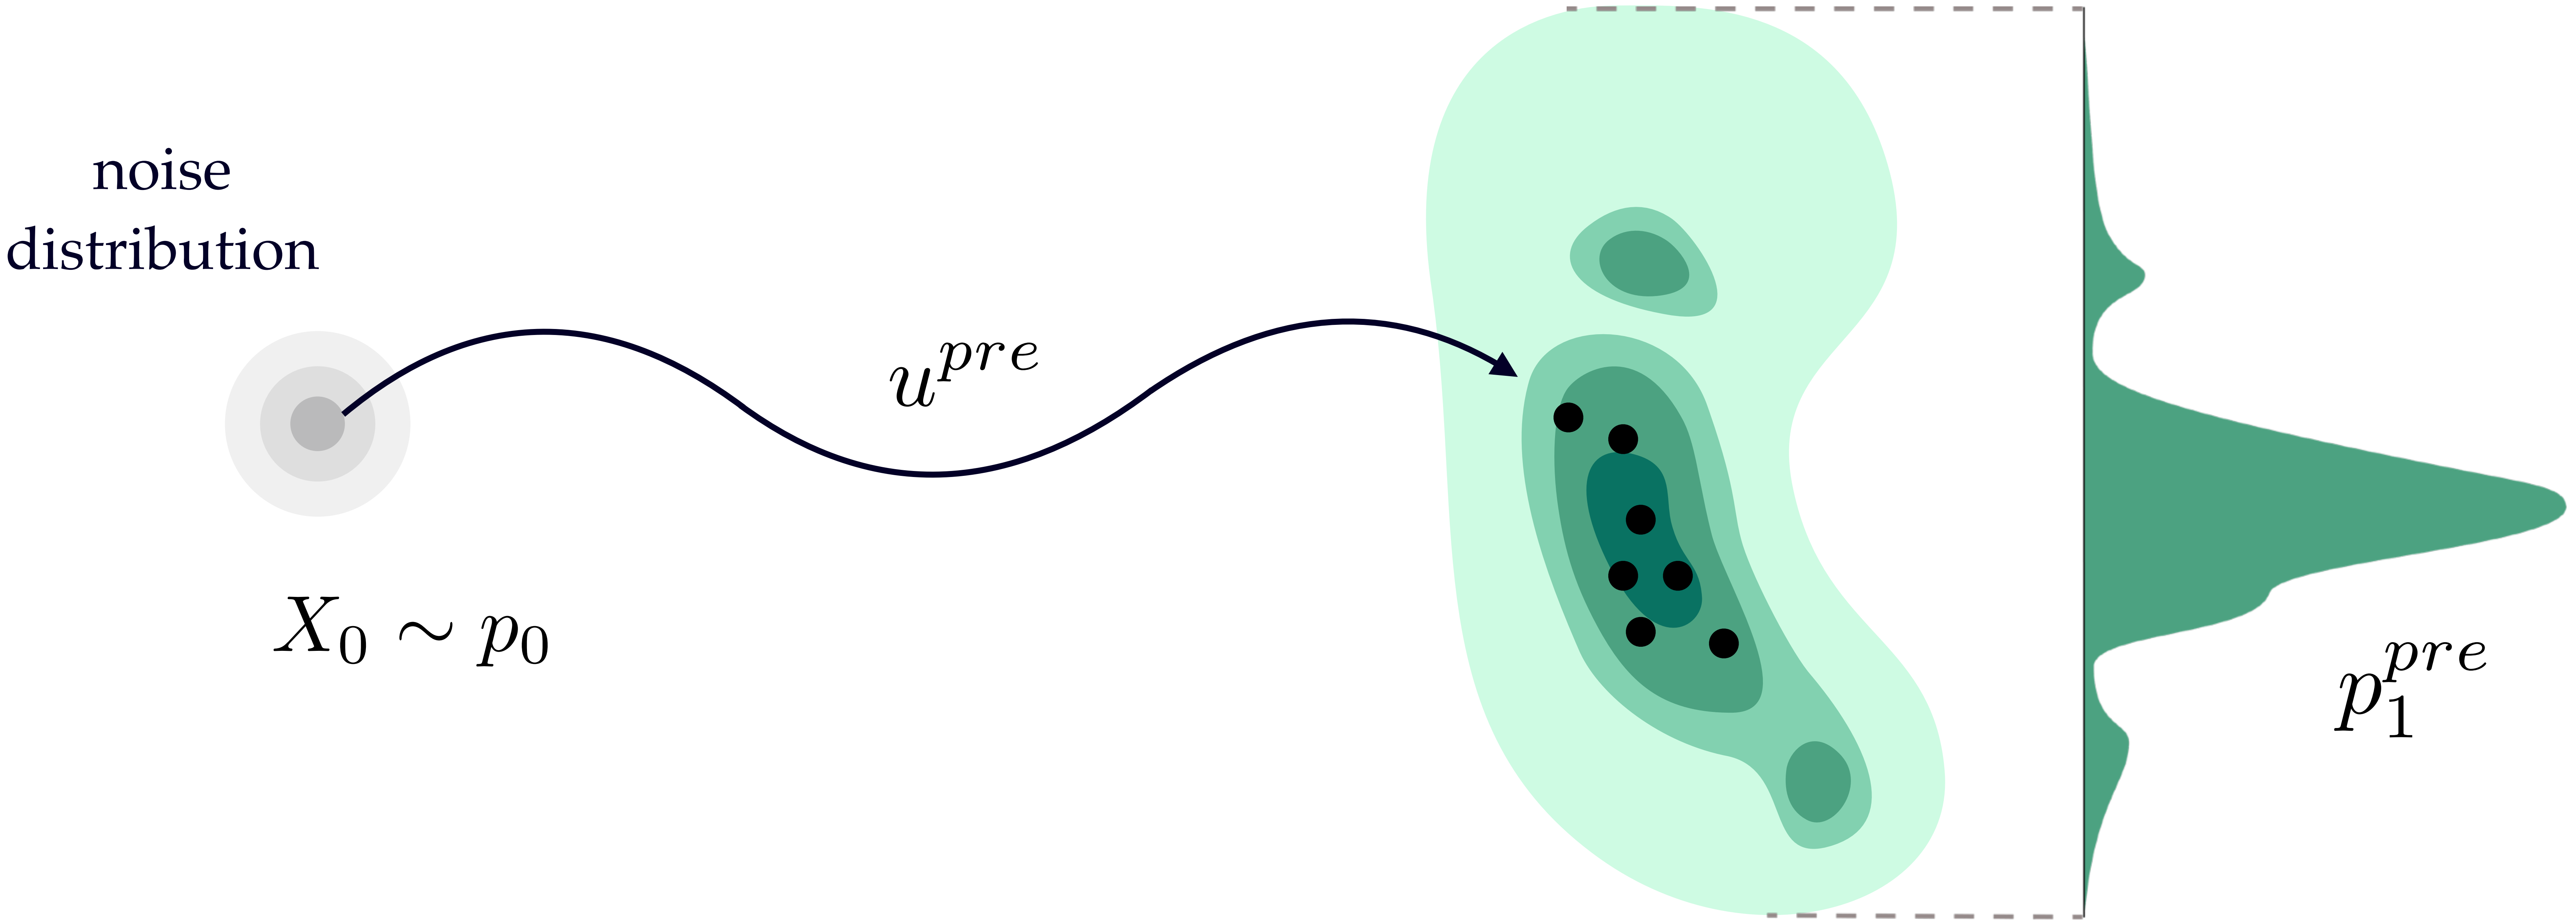

Generation via learned vector field

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

Reward-Guided Flow Adaptation

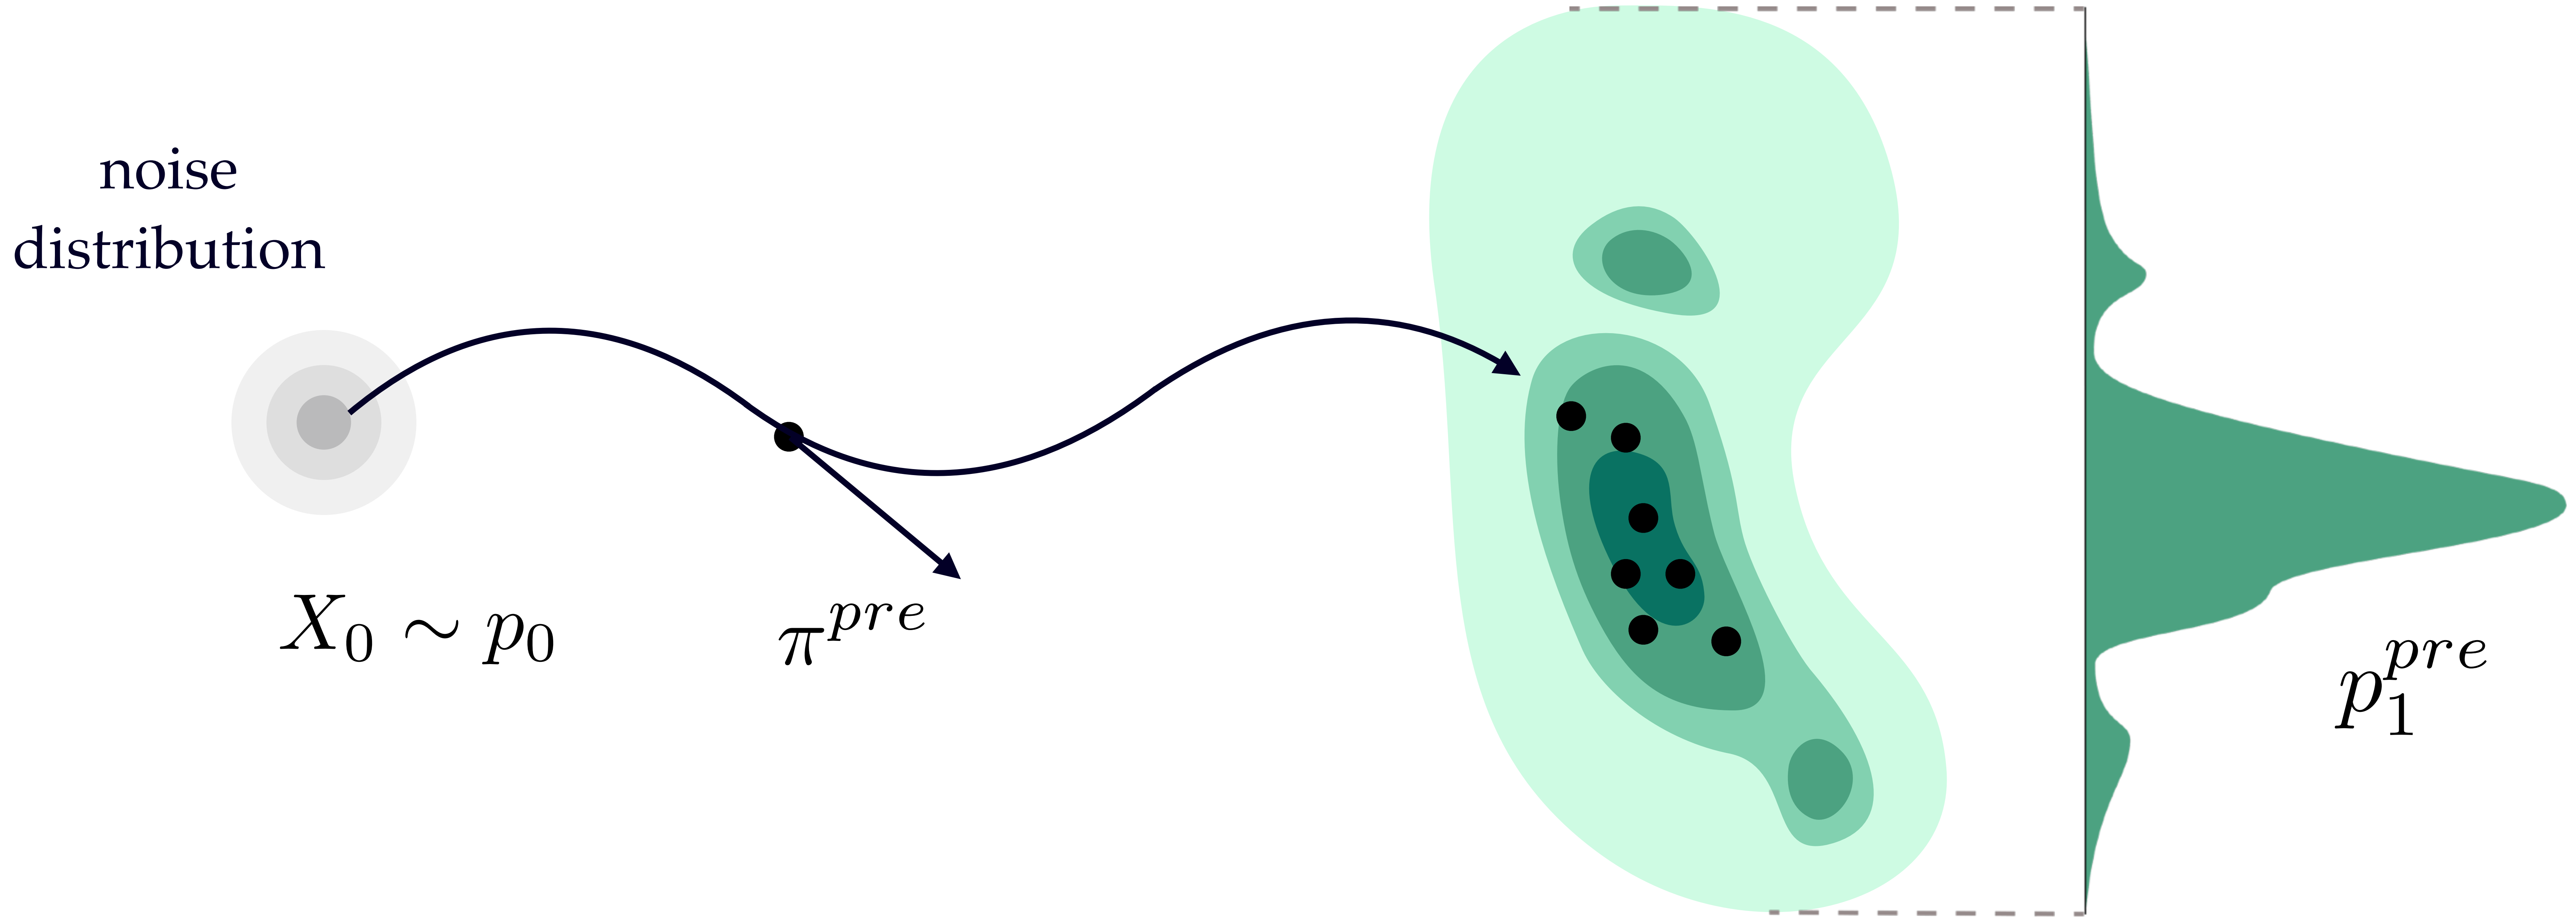

Velocity field as a **policy**

$$\pi^{pre}(X_t, t) := u^{pre}(X_t, t)$$

Reward-Guided Flow Adaptation

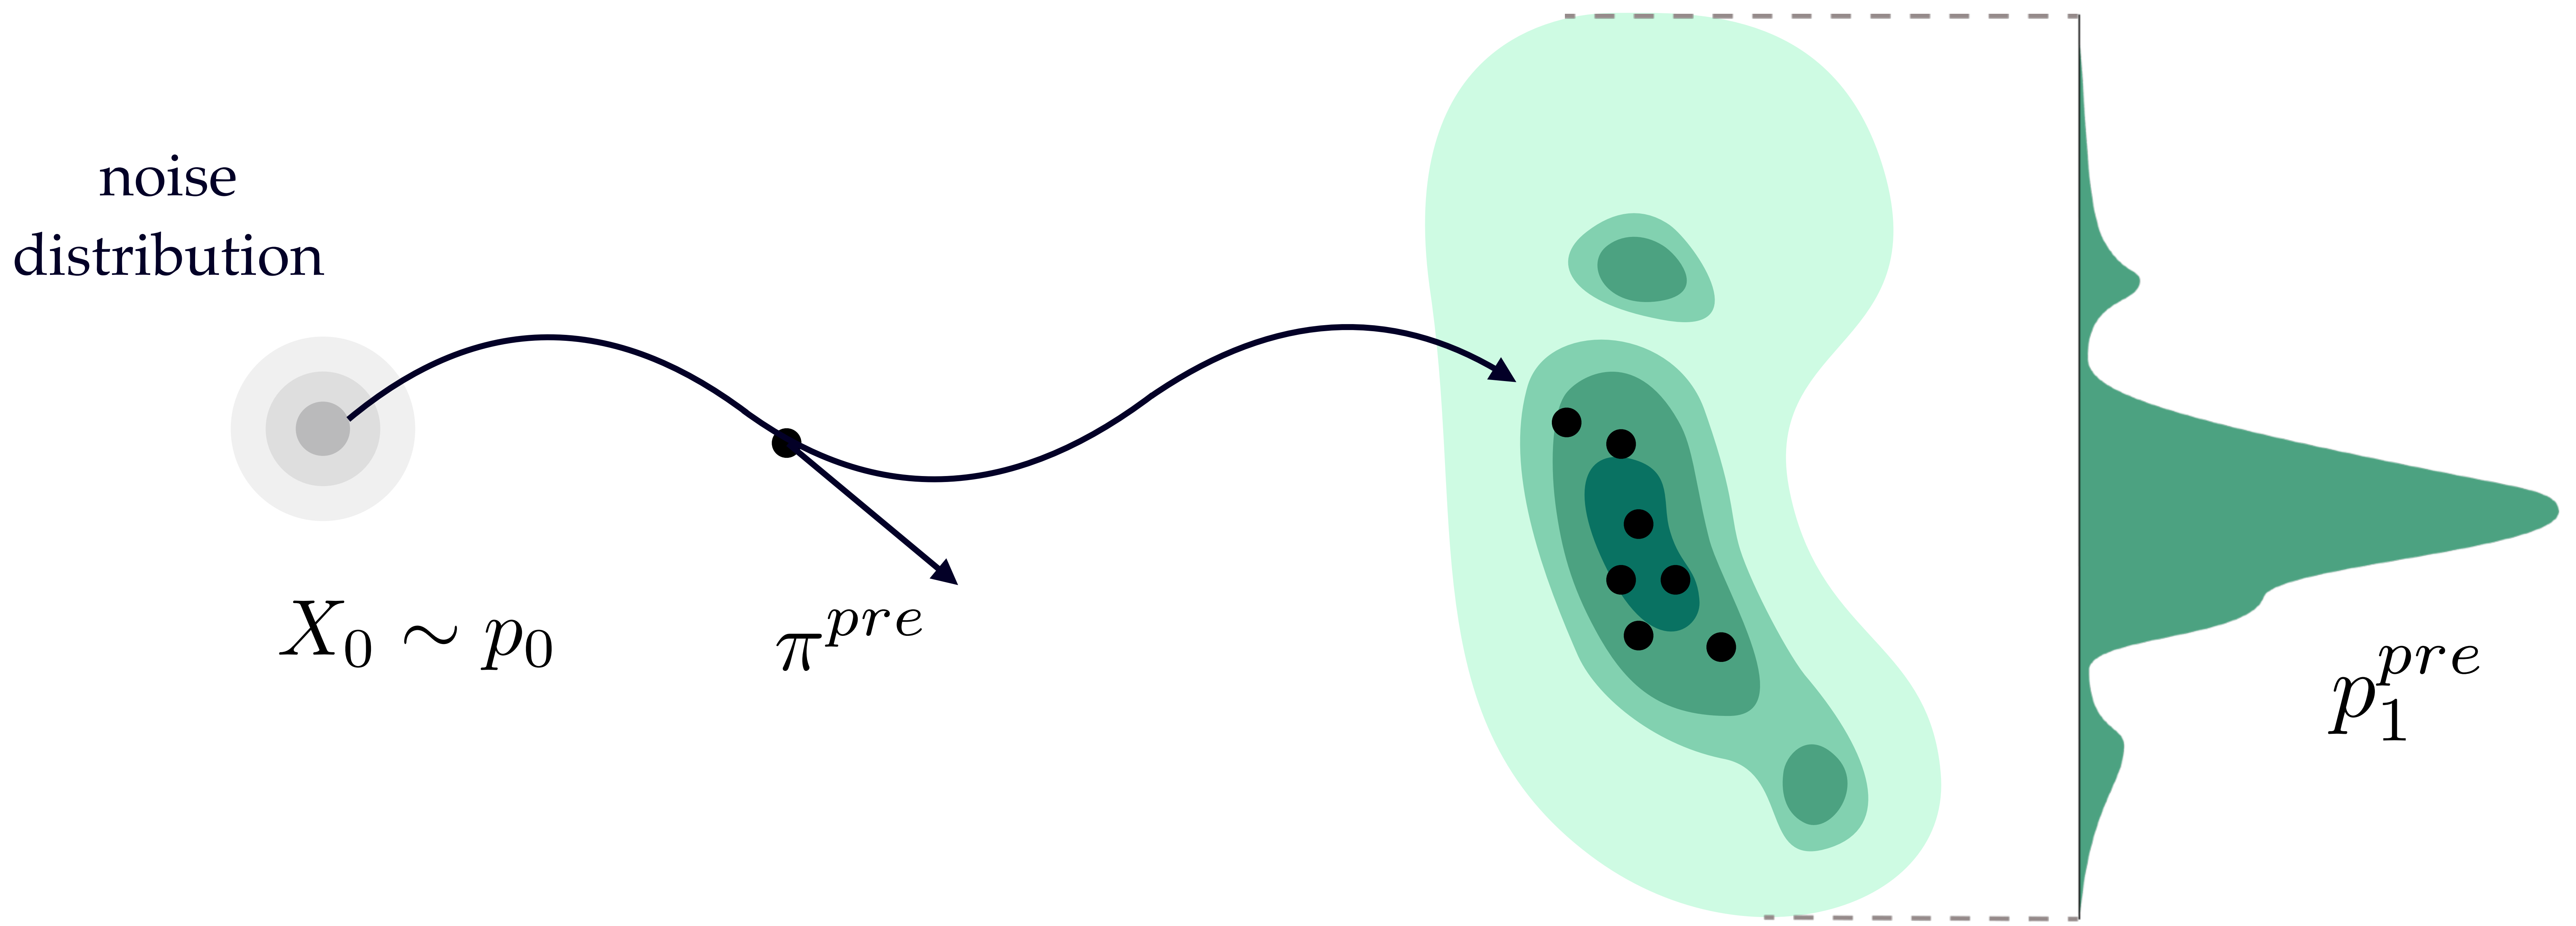

Velocity field as a **policy**

$$\pi^{pre}(X_t, t) := u^{pre}(X_t, t)$$

→

Marginal density induced by π^{pre}

$$p_1^{pre} := p_1^{\pi^{pre}} \approx p_{data}$$

Generative Optimization via Flow Adaptation

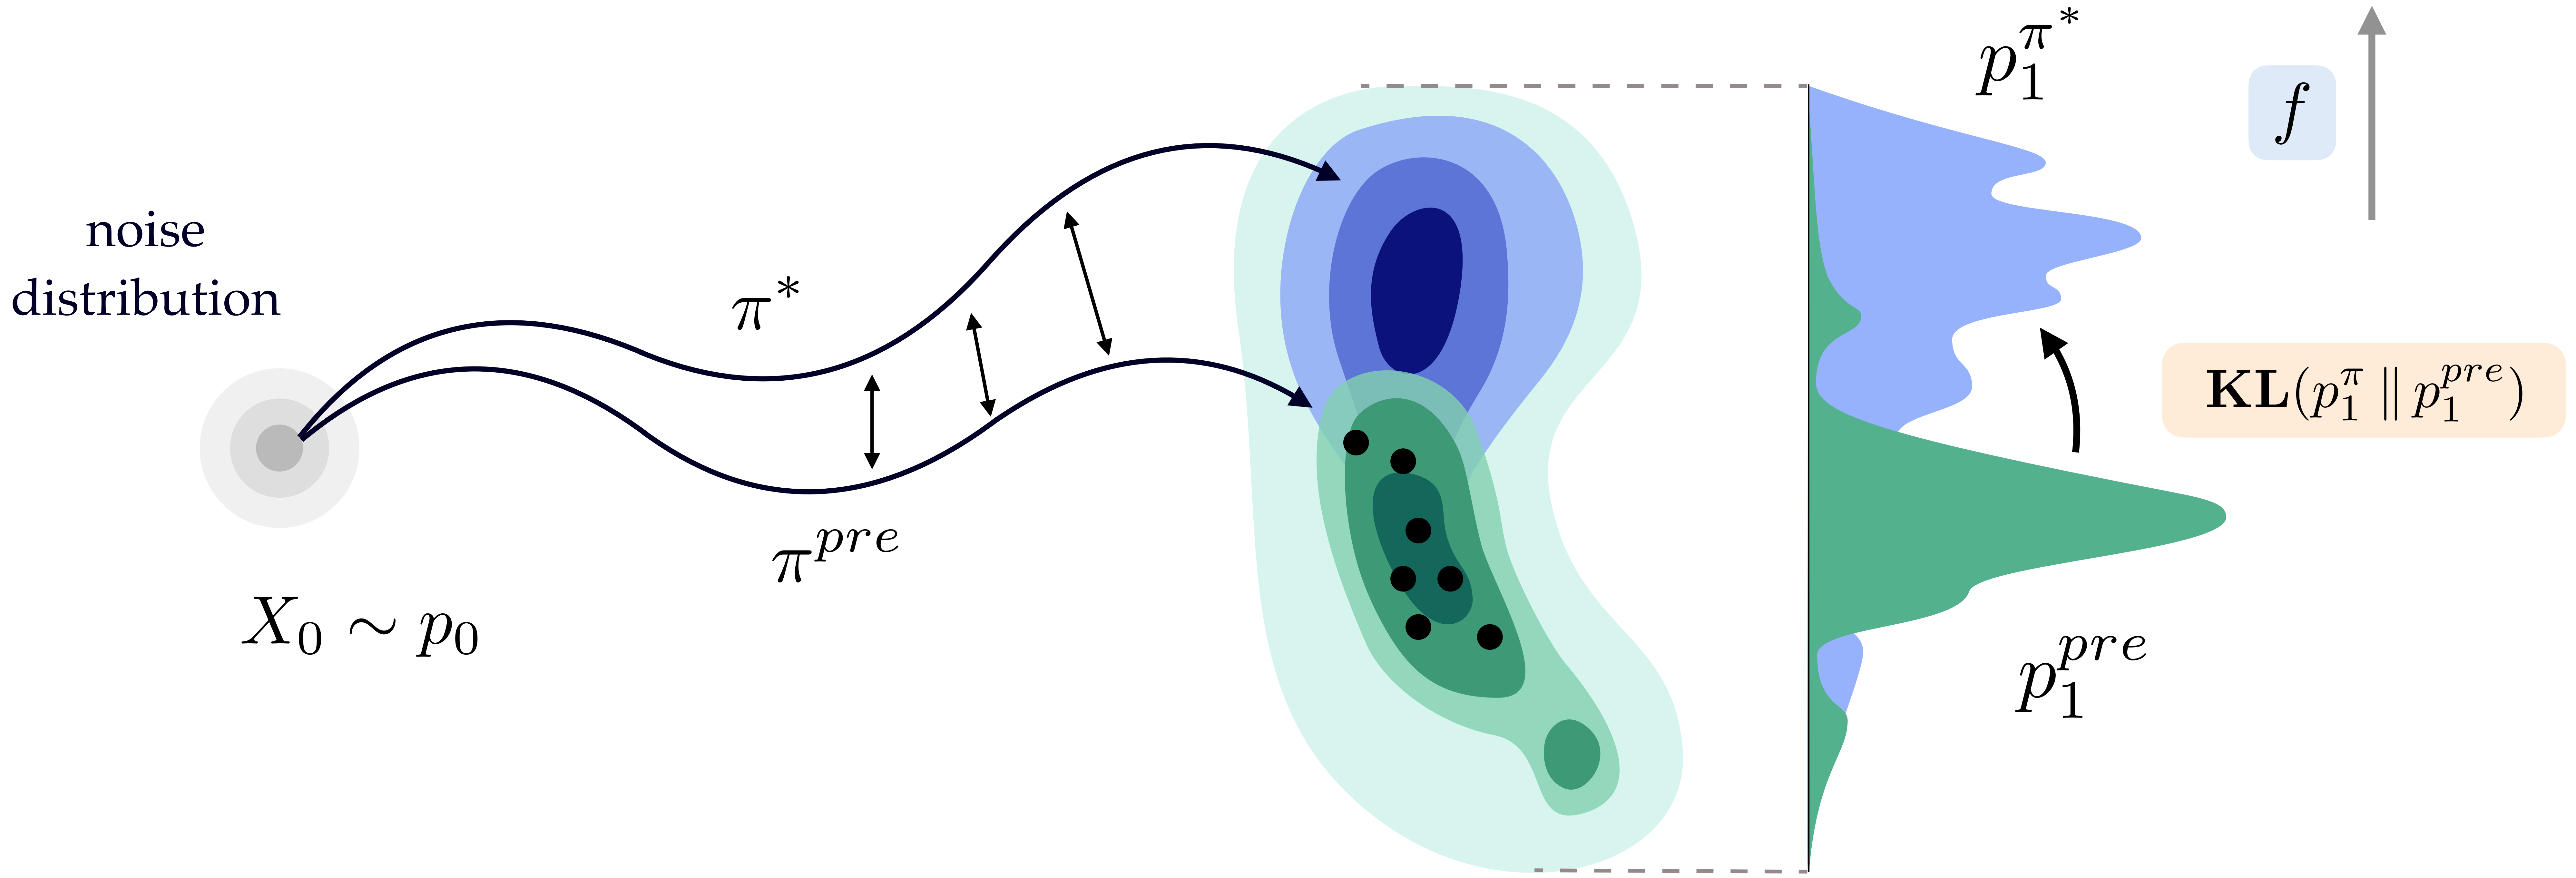

Flow Reward-Guided Adaptation

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

Generative Optimization via Flow Adaptation

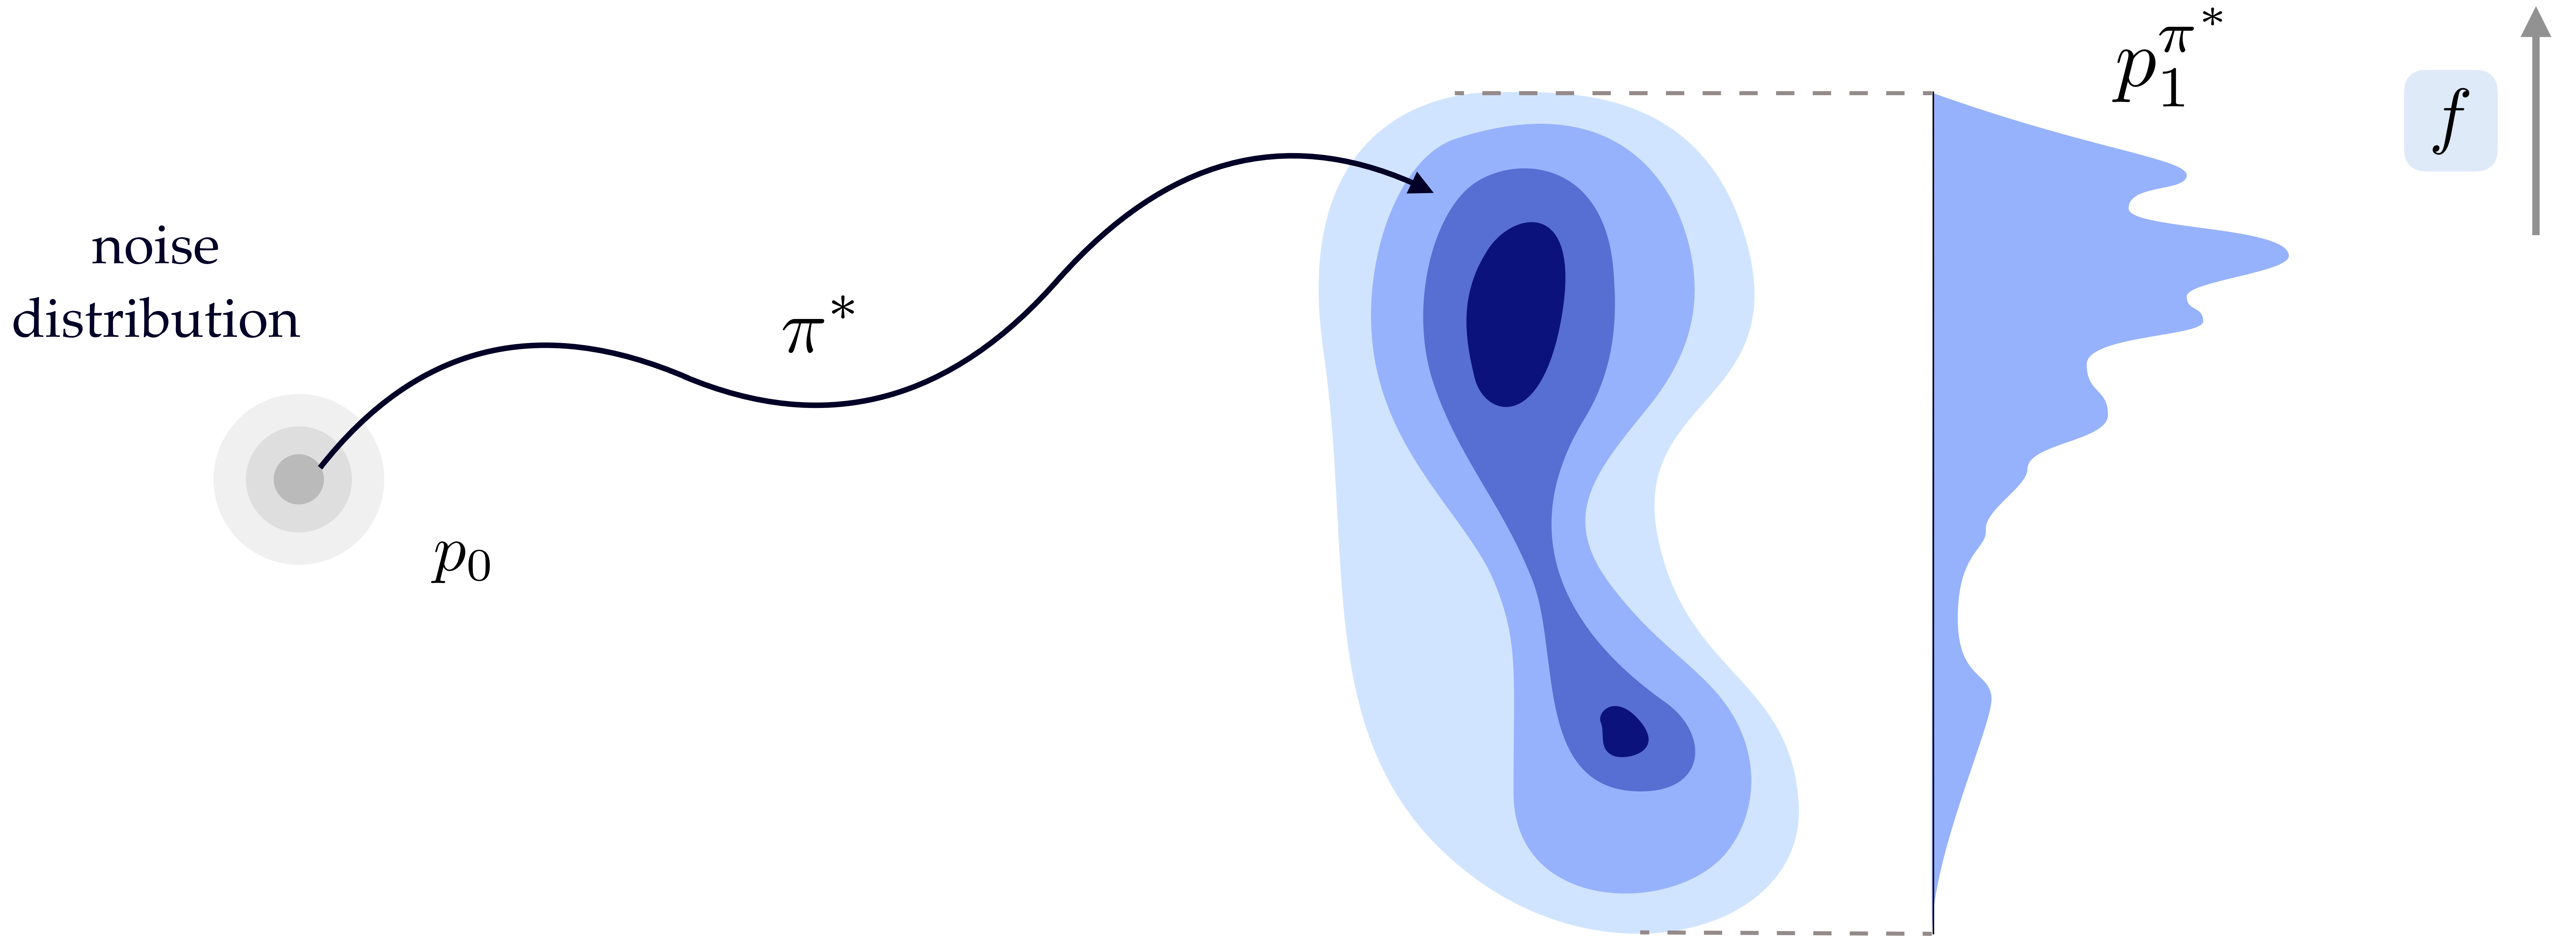

Flow Reward-Guided Adaptation

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

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

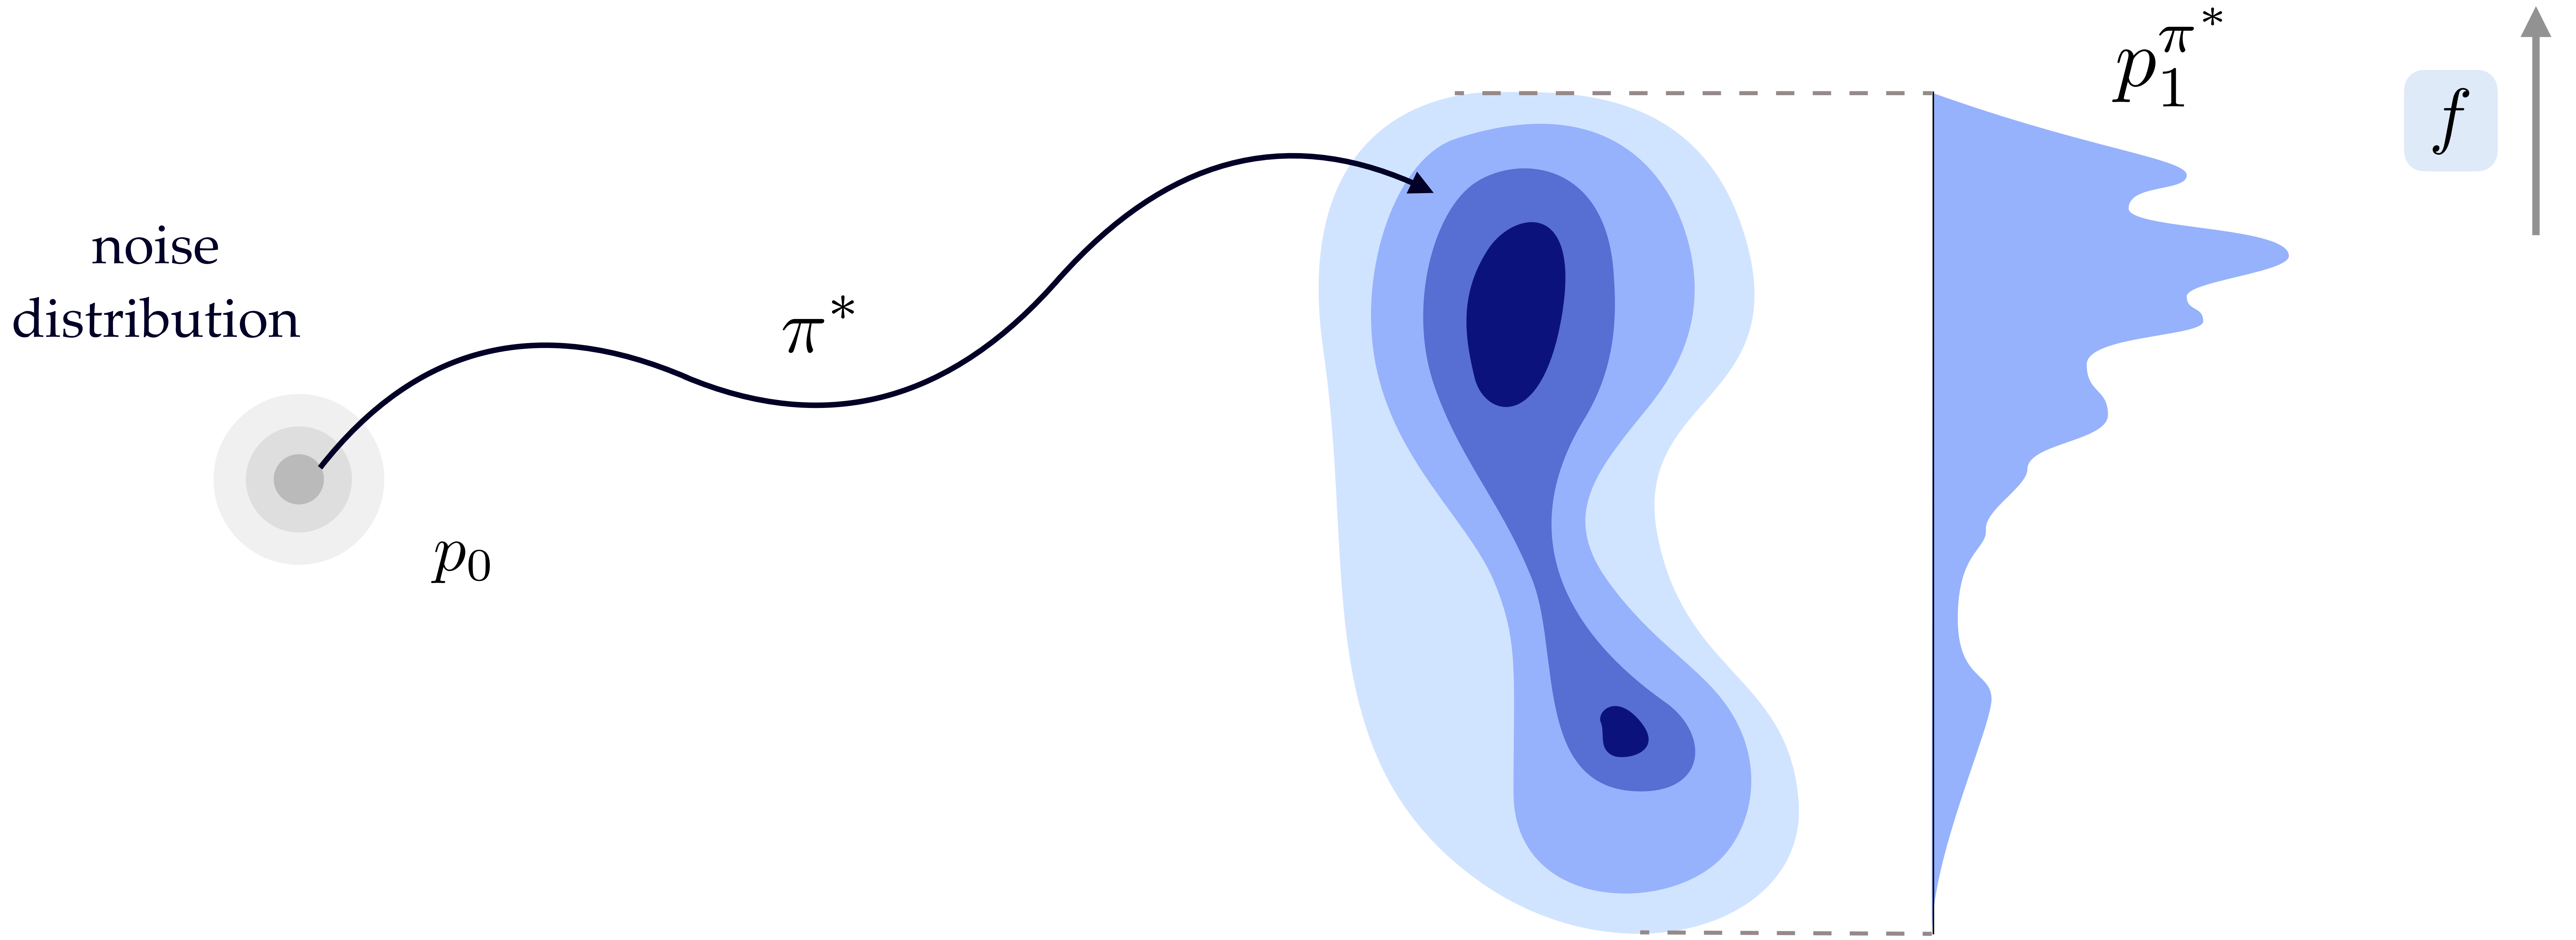

Flow Reward-Guided Adaptation

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

Example Methods.

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

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

Tail-Aware Reward Flow Adaptation: References

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

Is Expended Reward Maximization Good for Discovery?

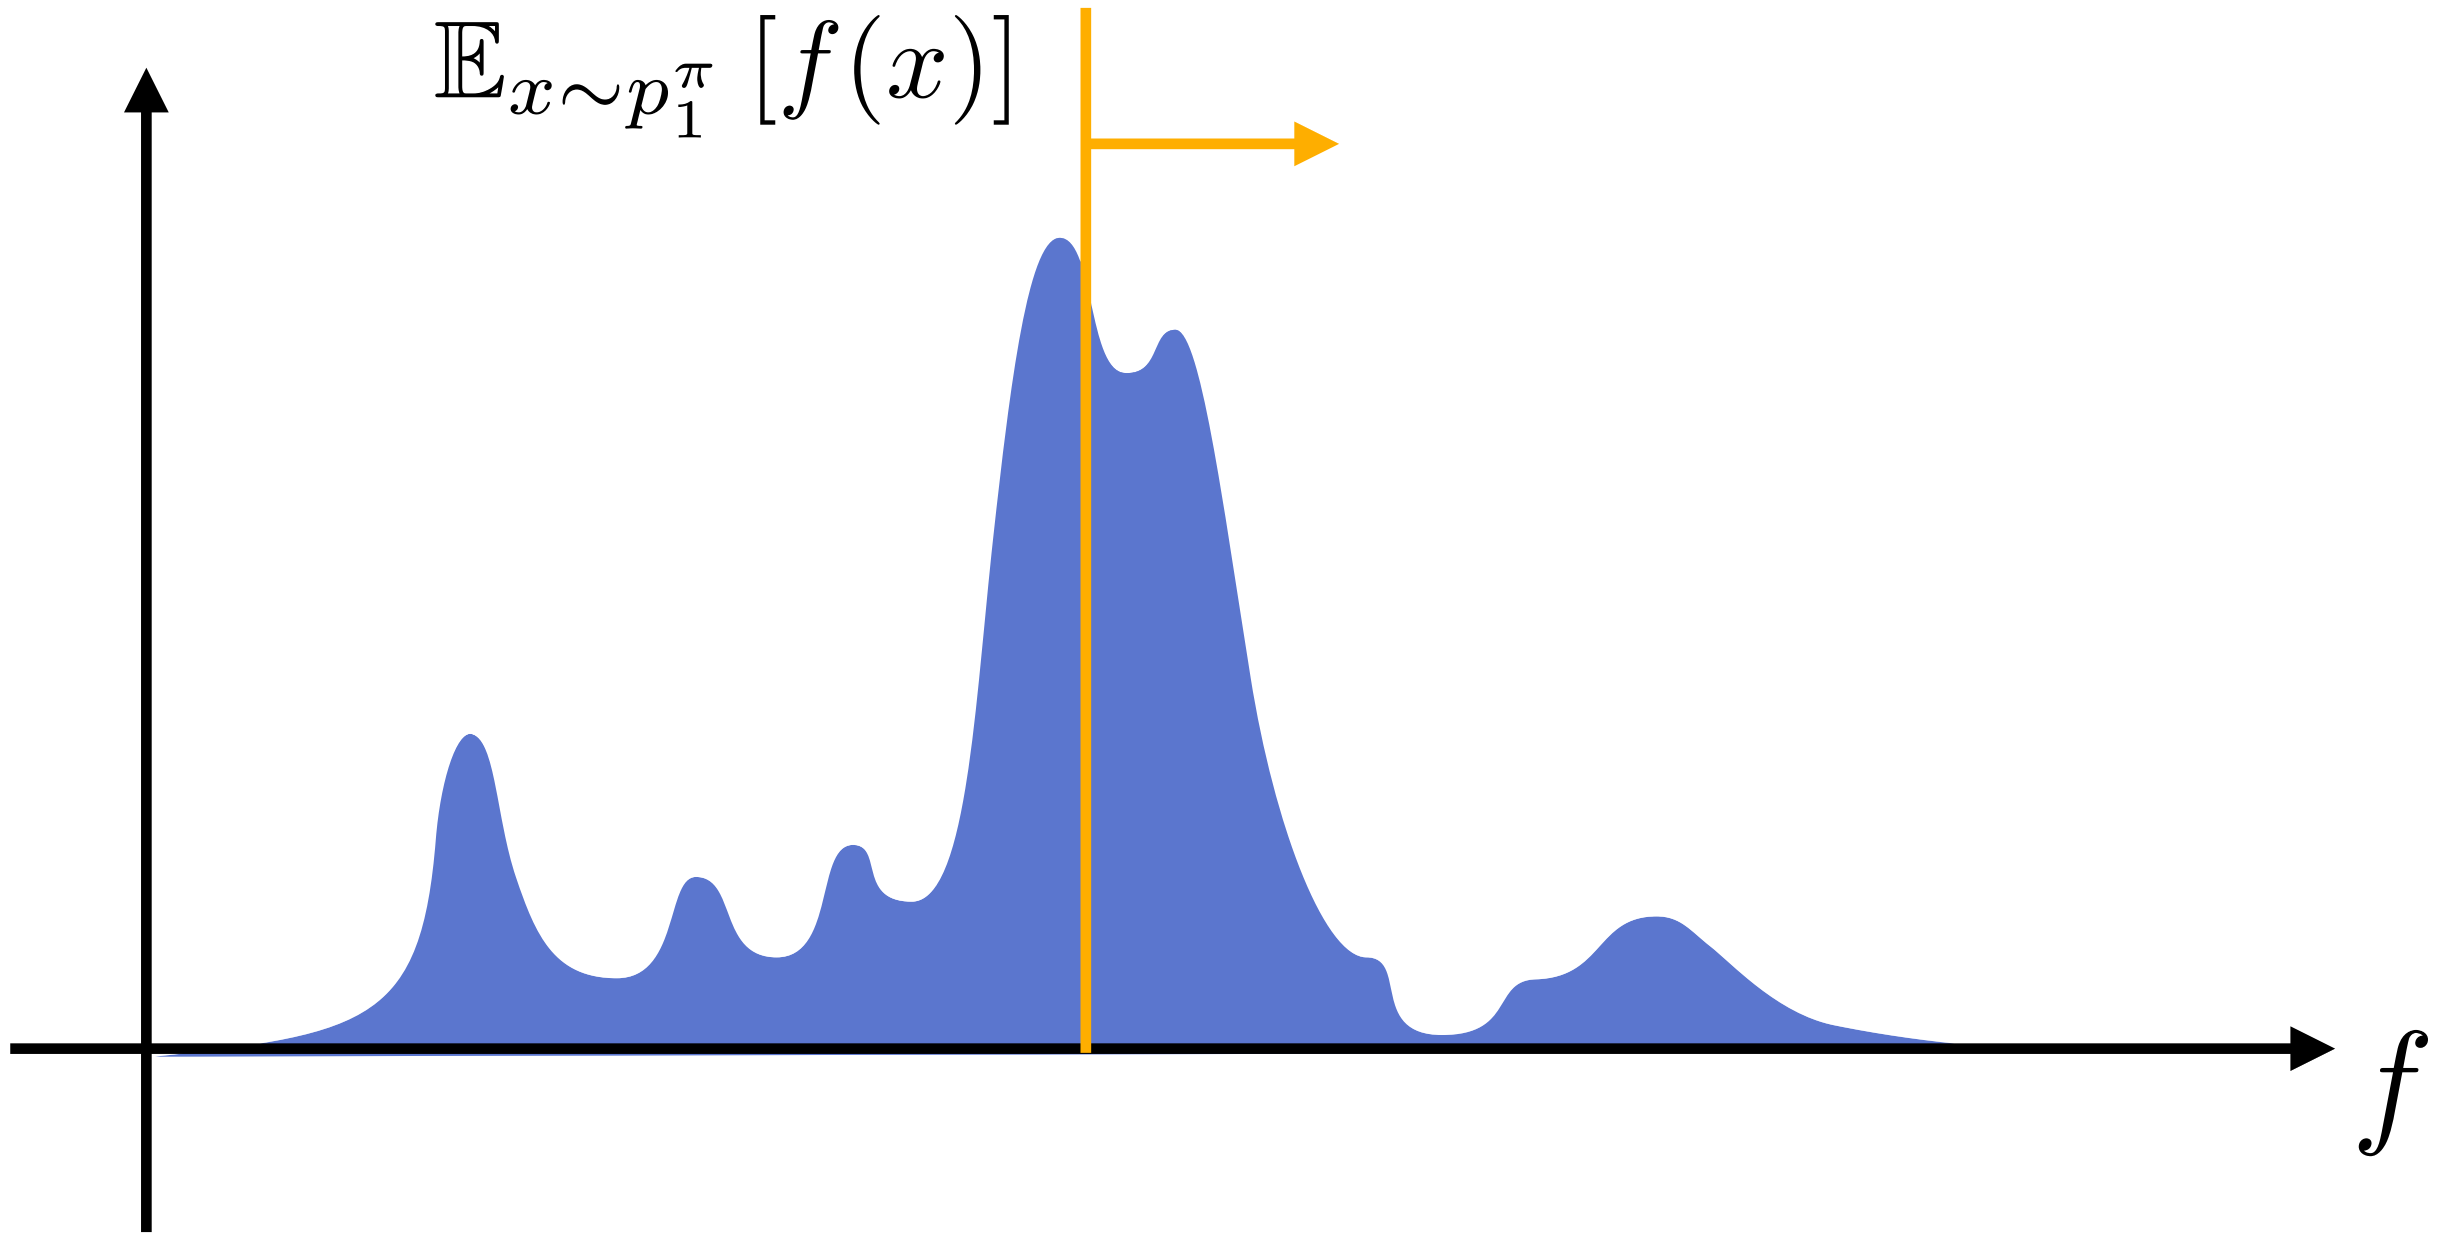

Flow Reward-Guided Fine-Tuning

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

Optimizes average sample quality

Discovery Pipeline in Chemistry and Biology

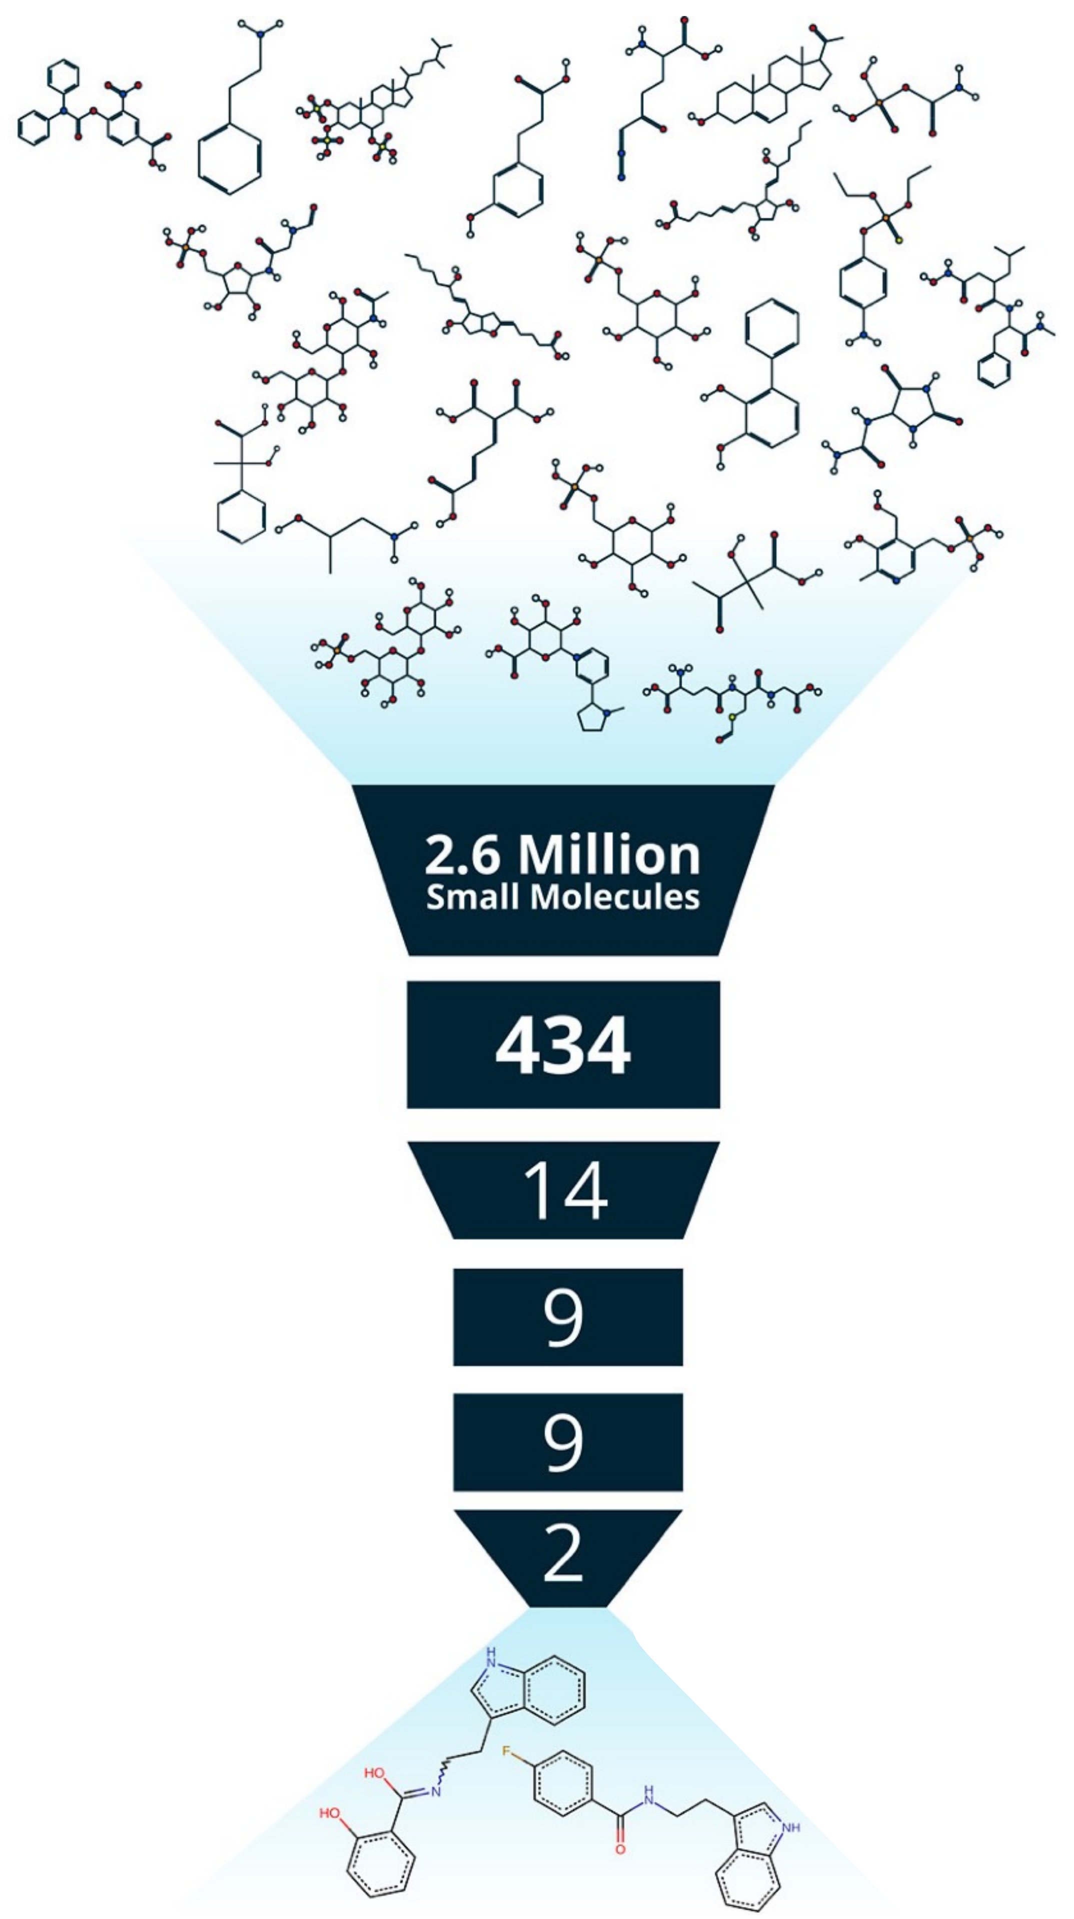

[Anastasiia Gryniukova et al., 2023]

Discovery Pipeline in Chemistry and Biology

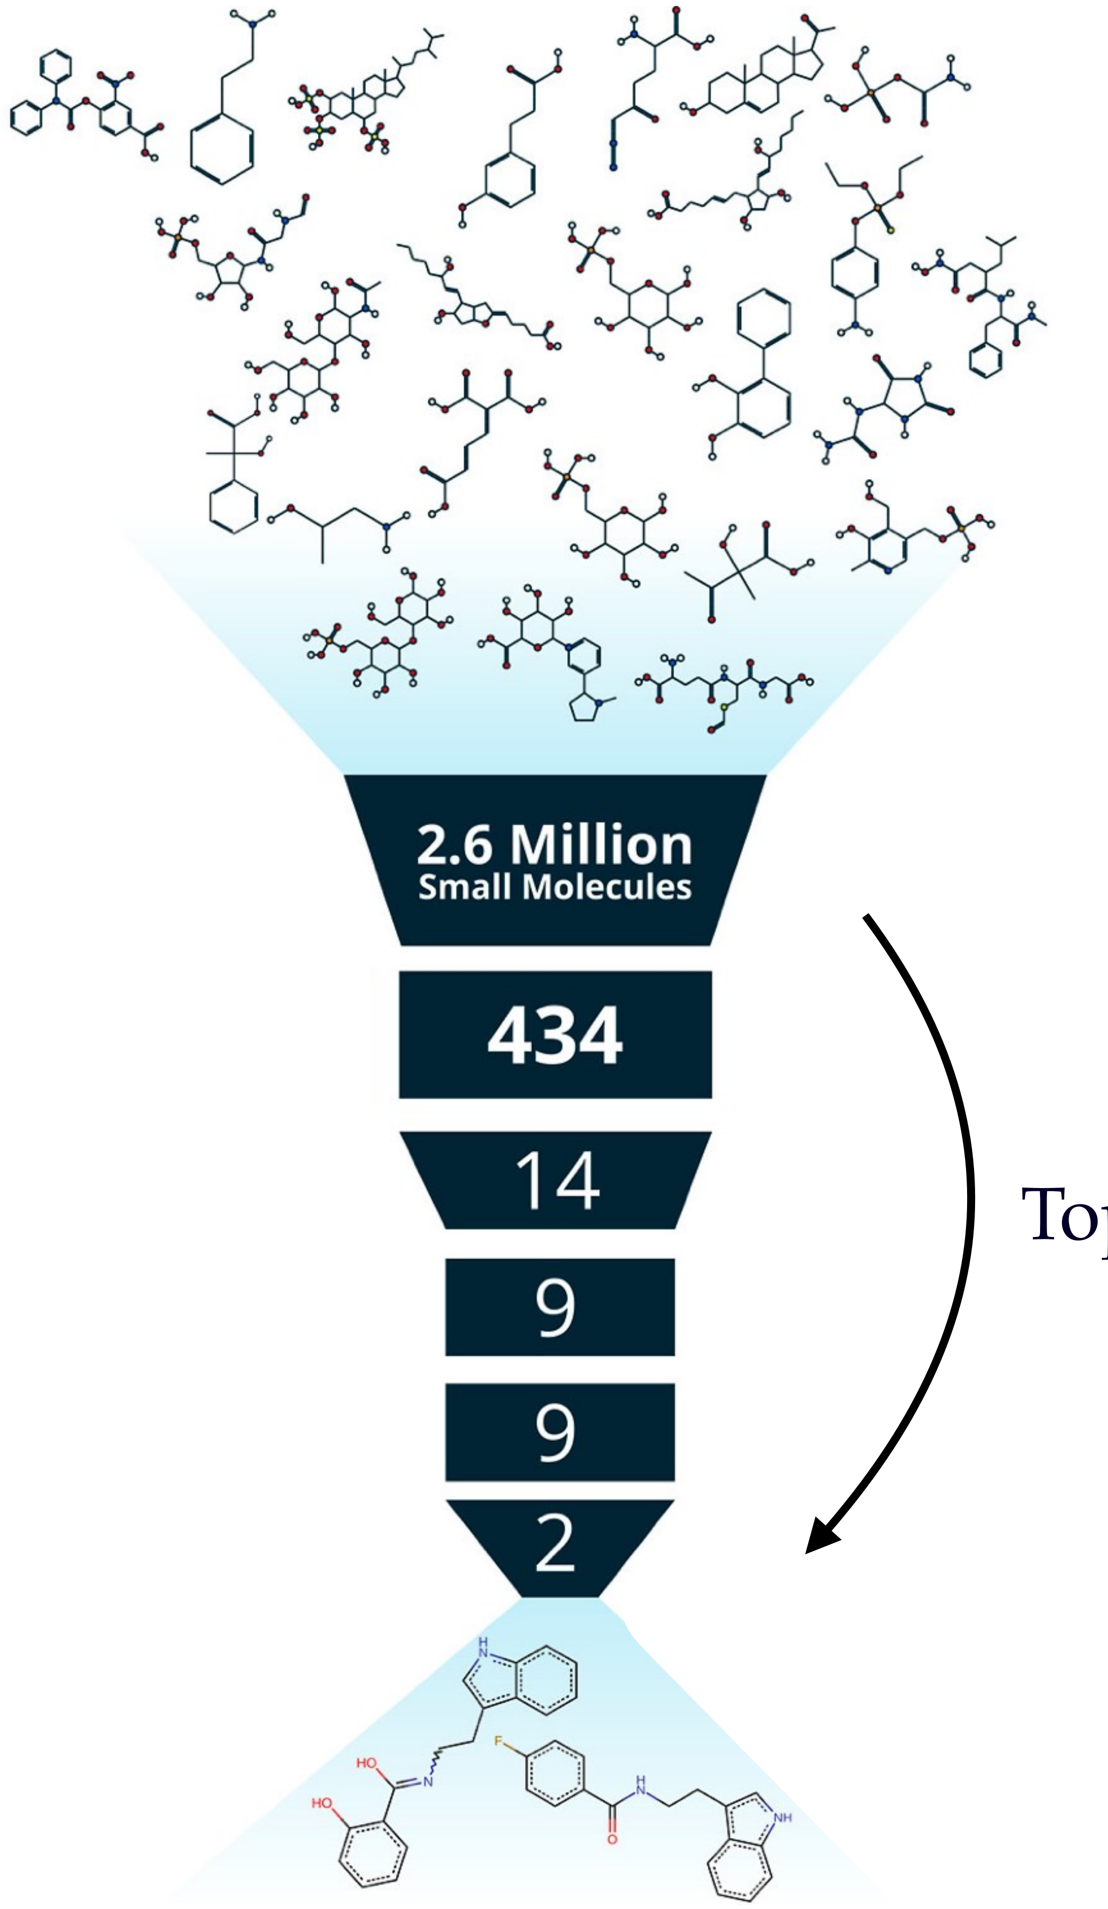

Top $0.00...x$ % molecules

[Anastasiia Gryniukova et al., 2023]

Discovery Pipeline in Chemistry and Biology

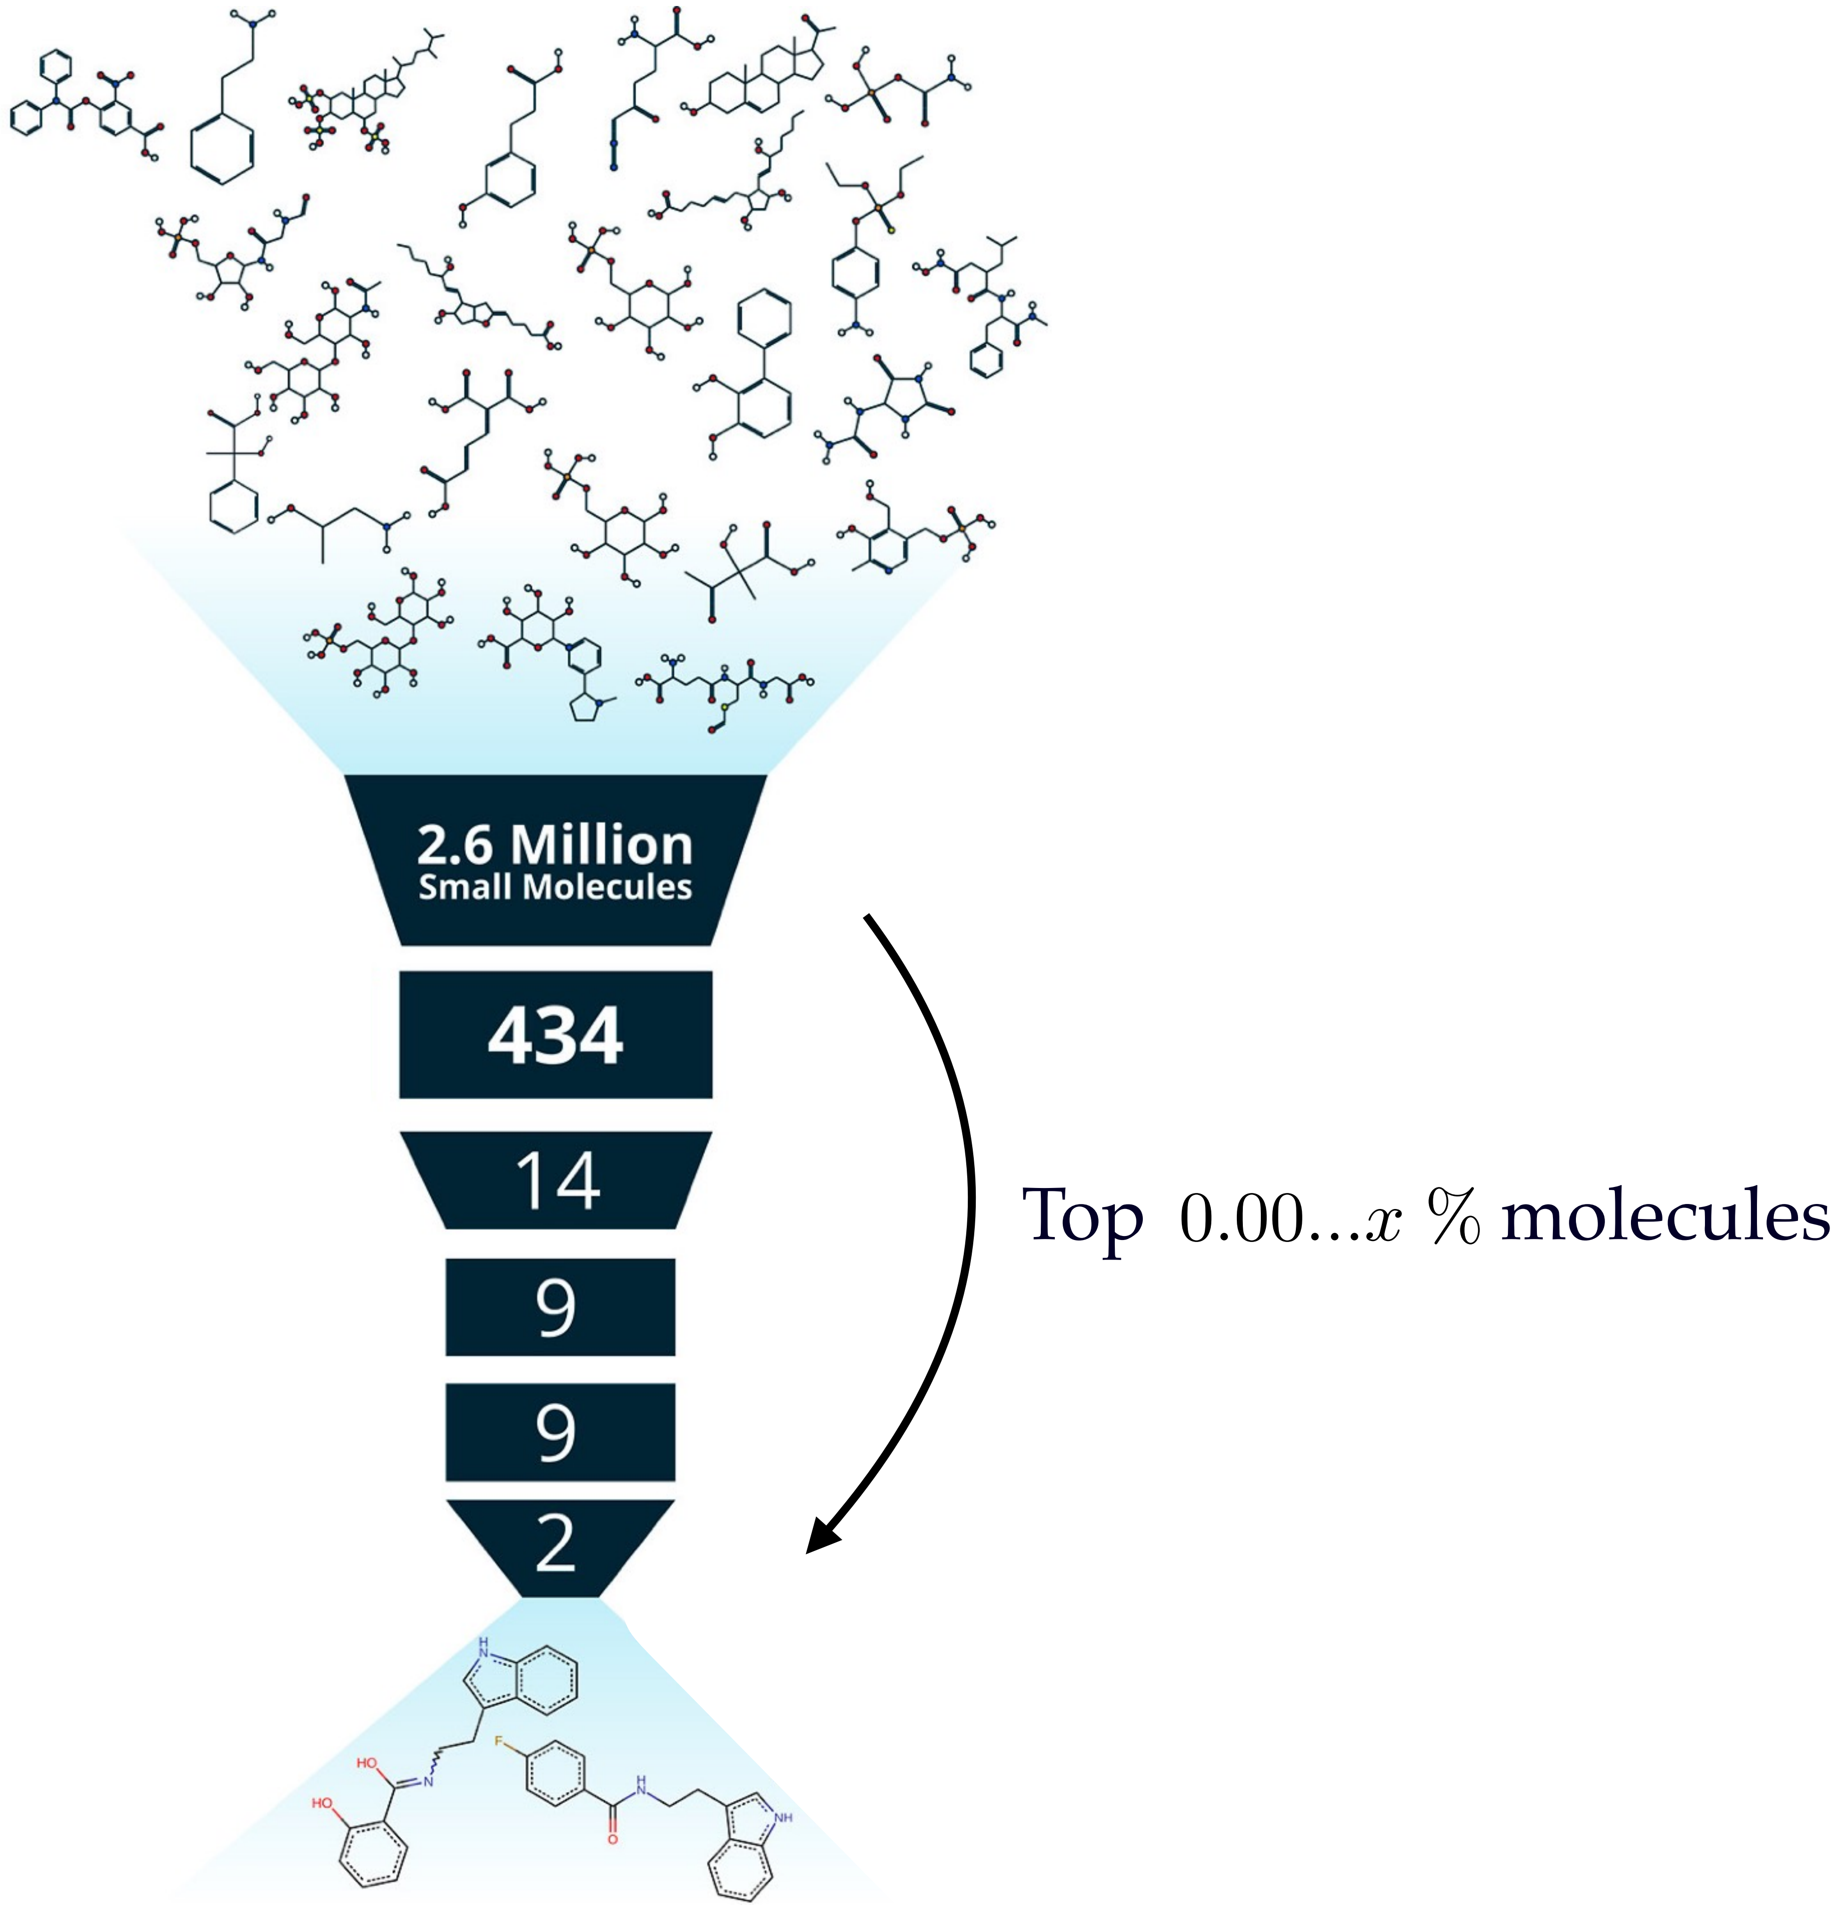

[Anastasiia Gryniukova et al., 2023]

Takeaway

Average reward adaptation is often a bad objective for discovery.

Risk-sensitivity for Quantitative Finance

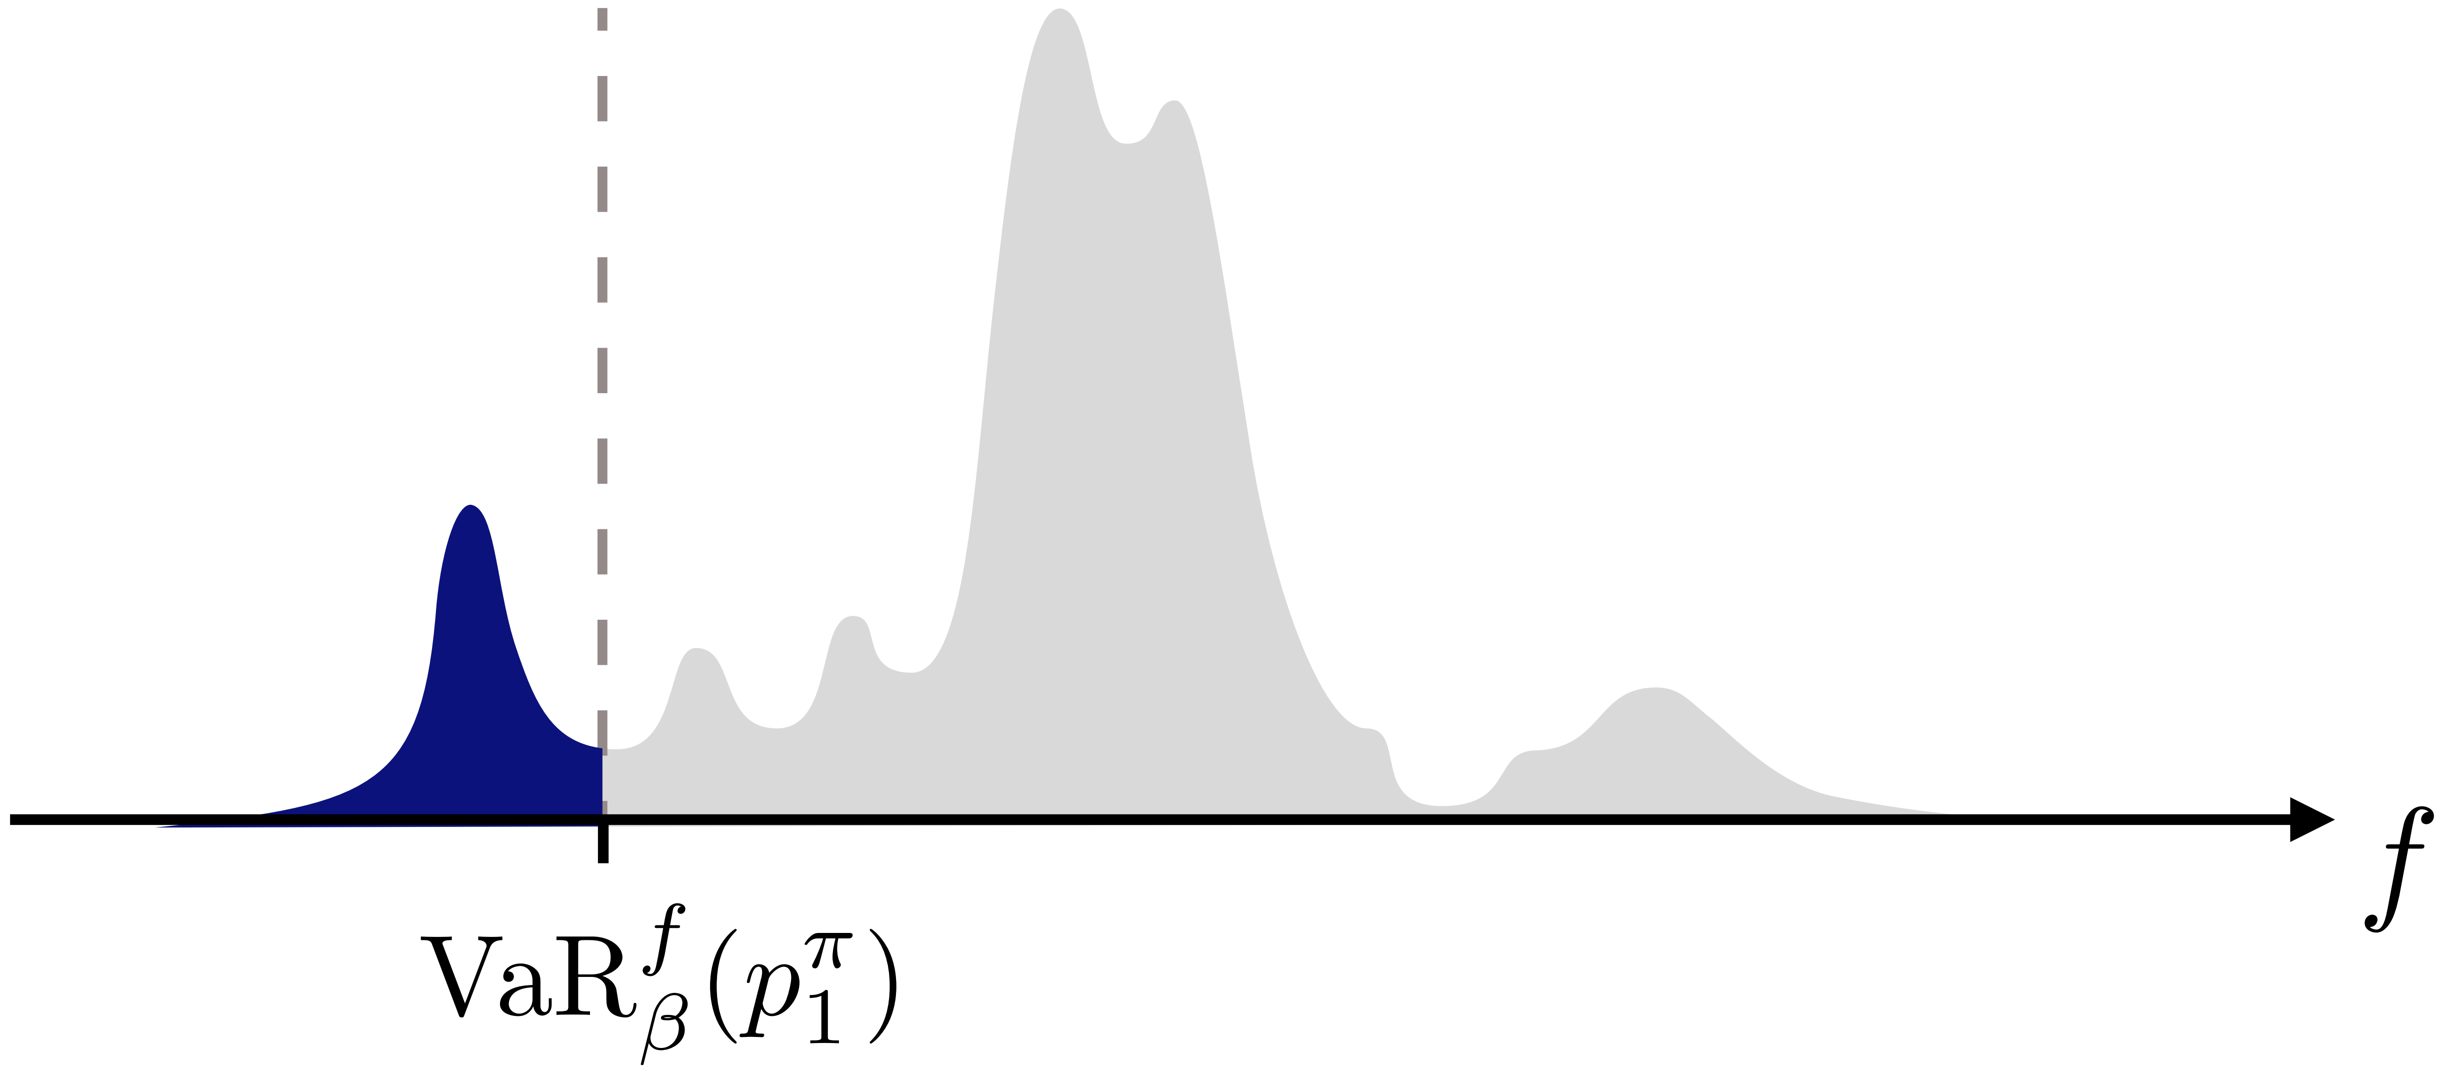

Risk-sensitivity for Quantitative Finance

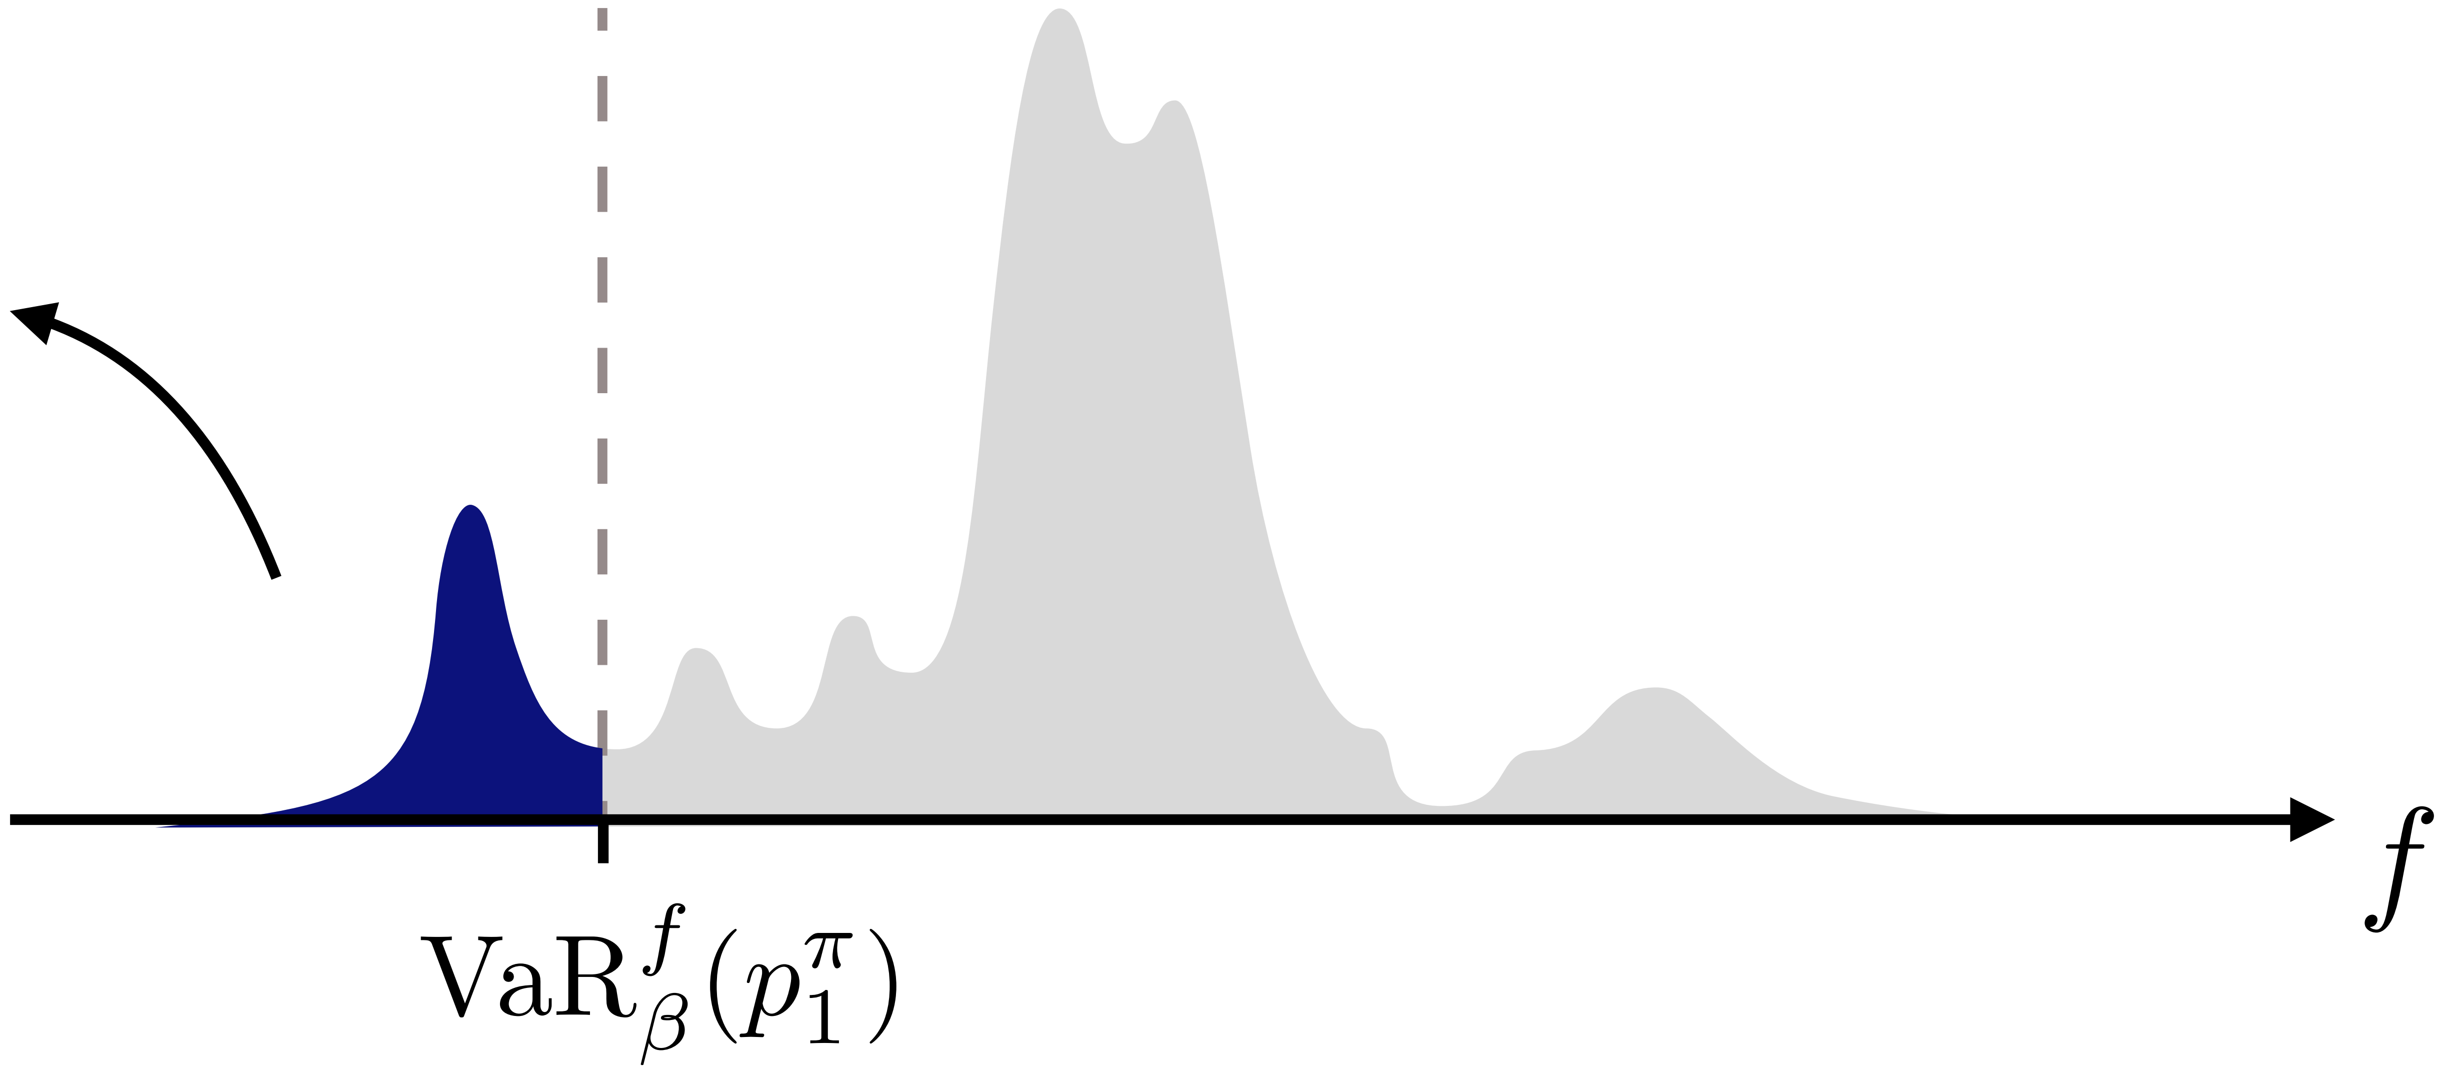

Risk-sensitivity for Quantitative Finance

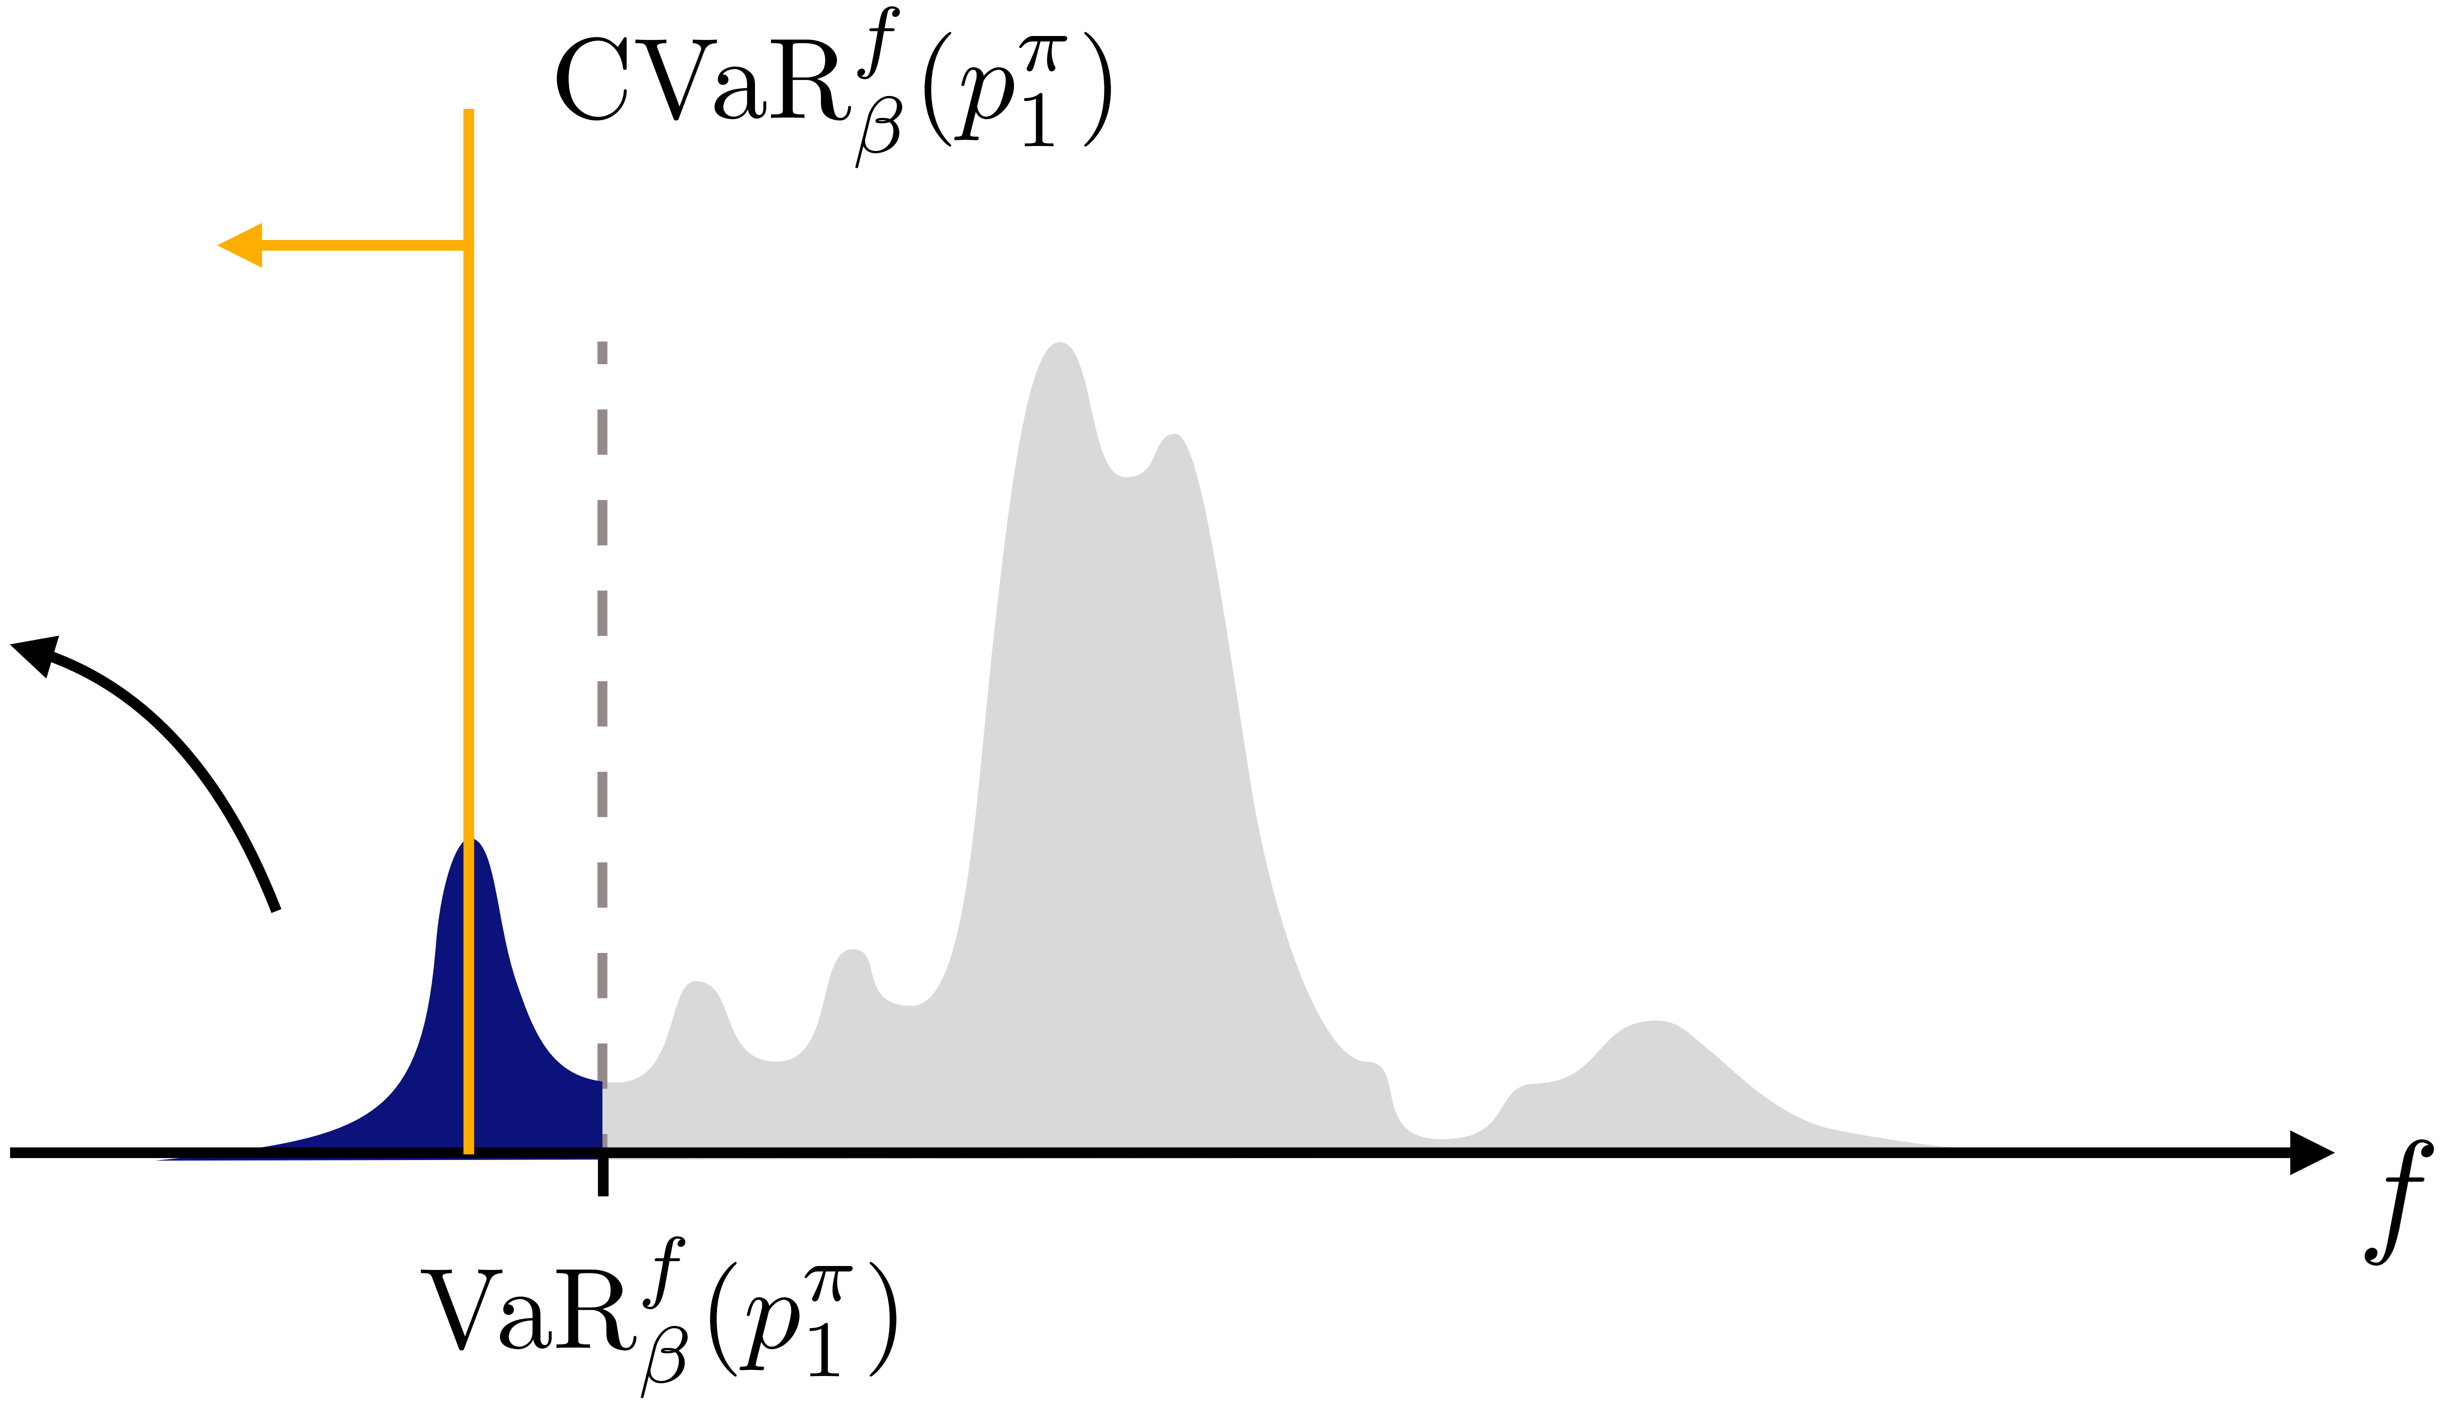

$$\mathrm{CVaR}_\beta^f(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \leq \mathrm{VaR}_\beta^f(p_1^\pi)]$$

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

Risk-Sensitive Flow Steering for Discovery

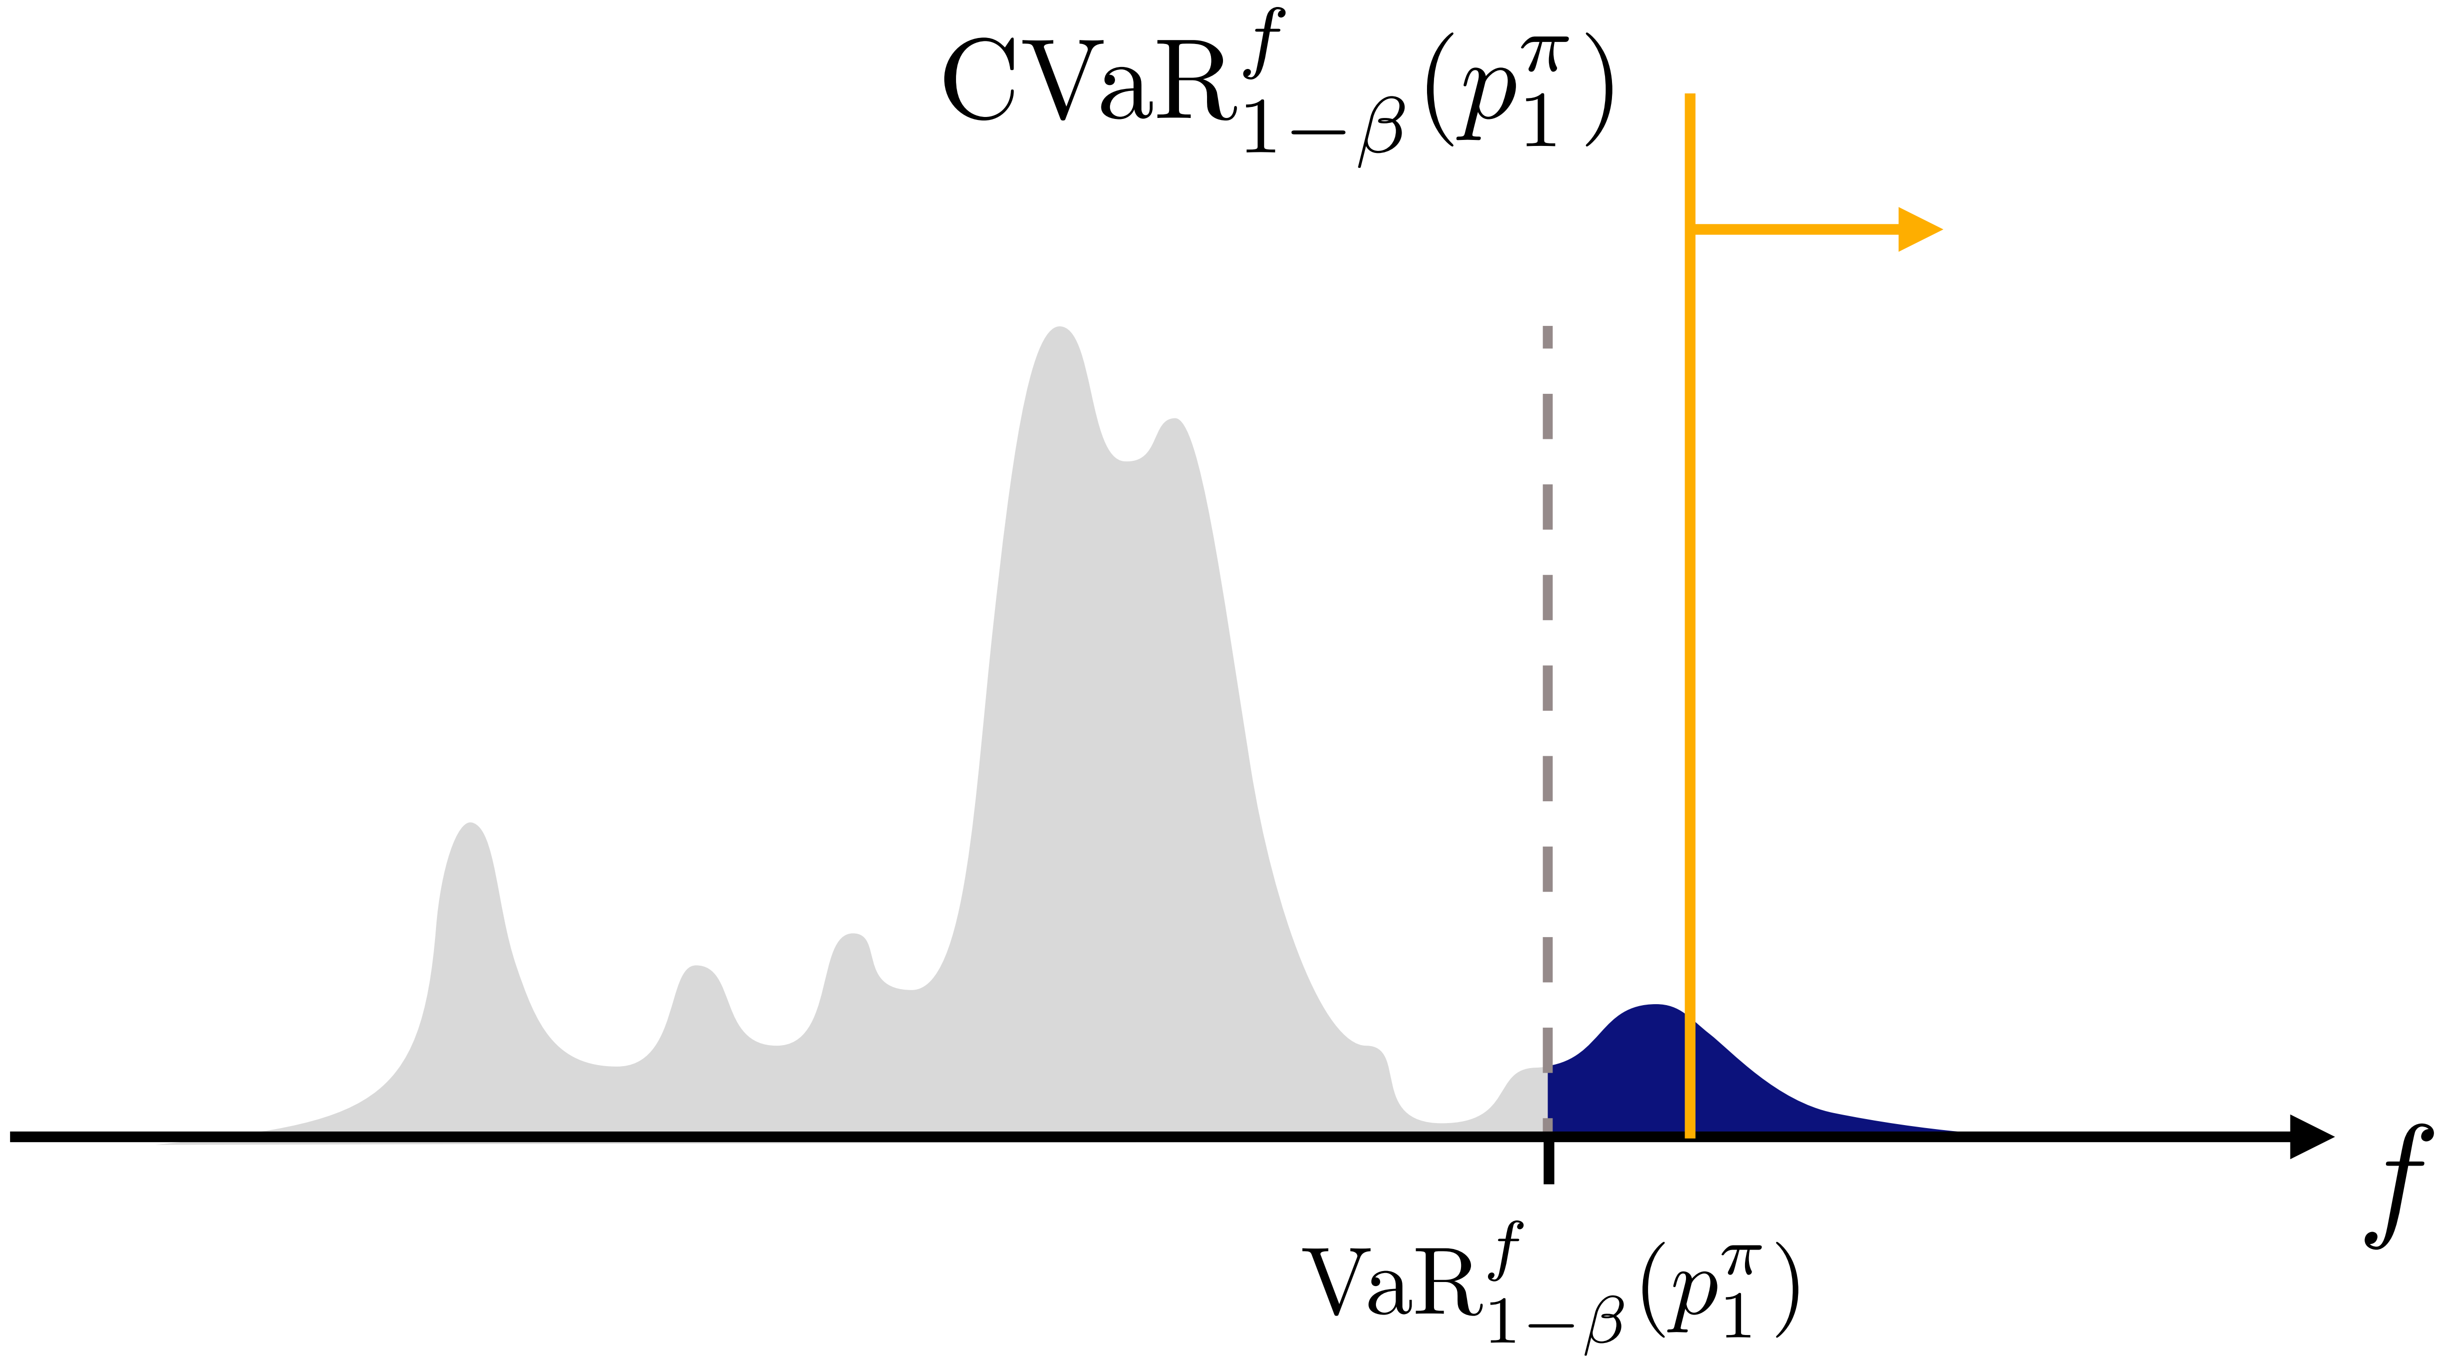

Risk-Sensitive Flow Steering for Discovery

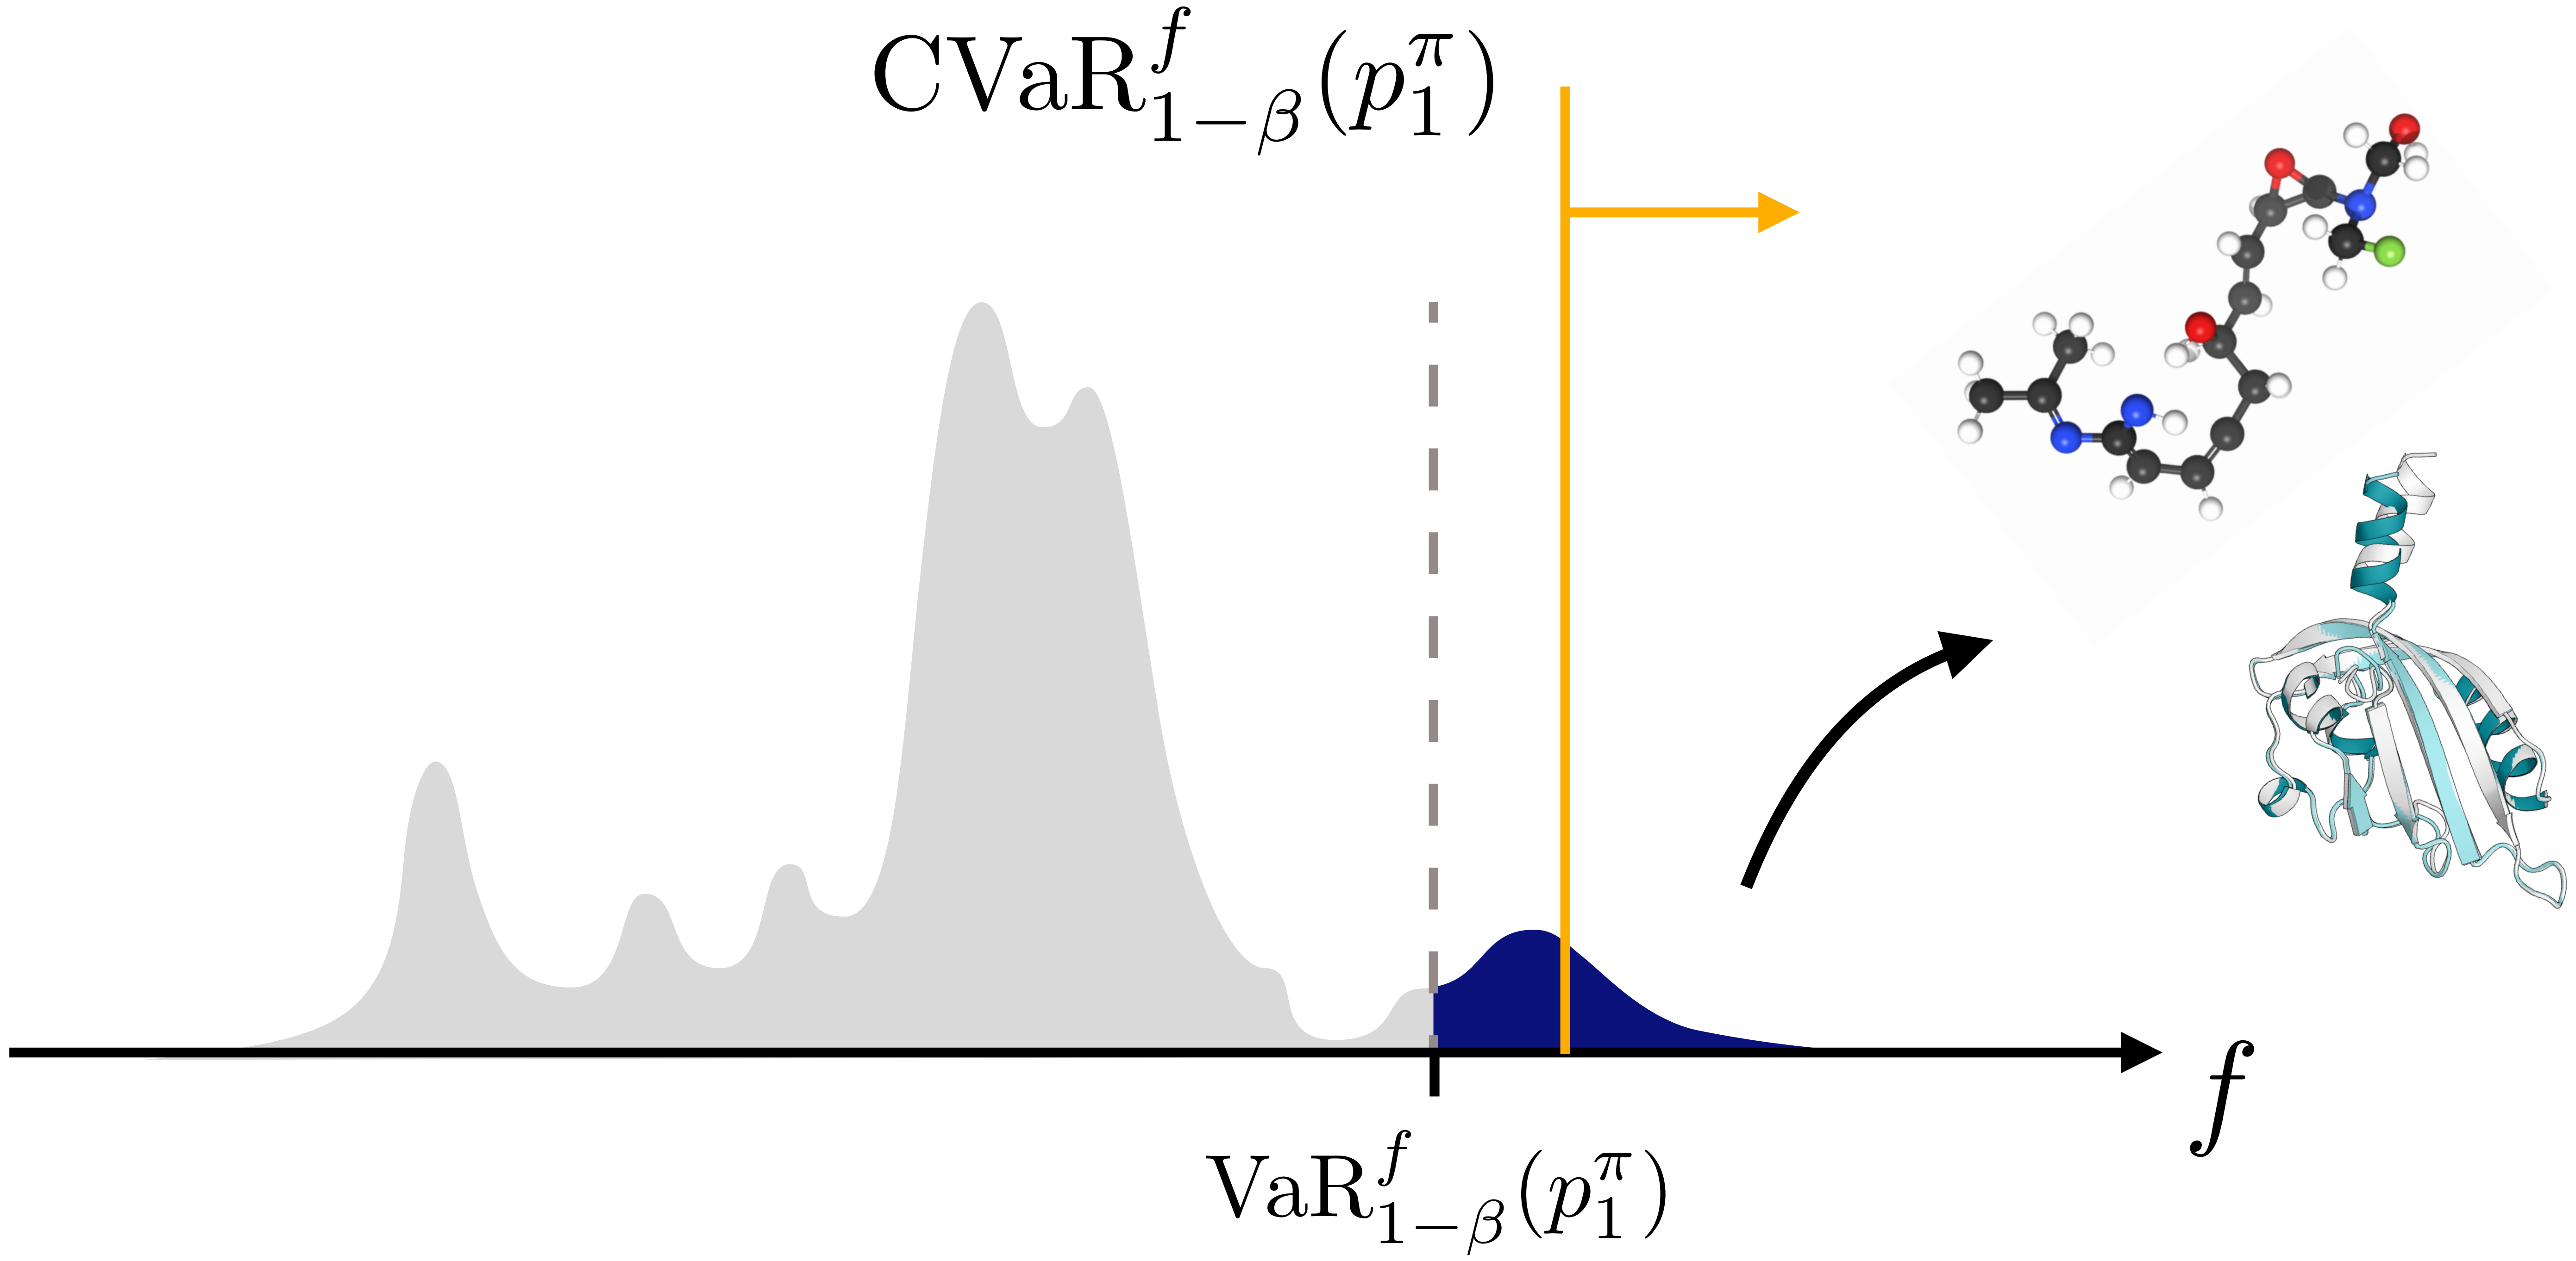

Valid high-value rare events $\approx$ (potential) discoveries

Risk-Sensitive Flow Steering for Discovery

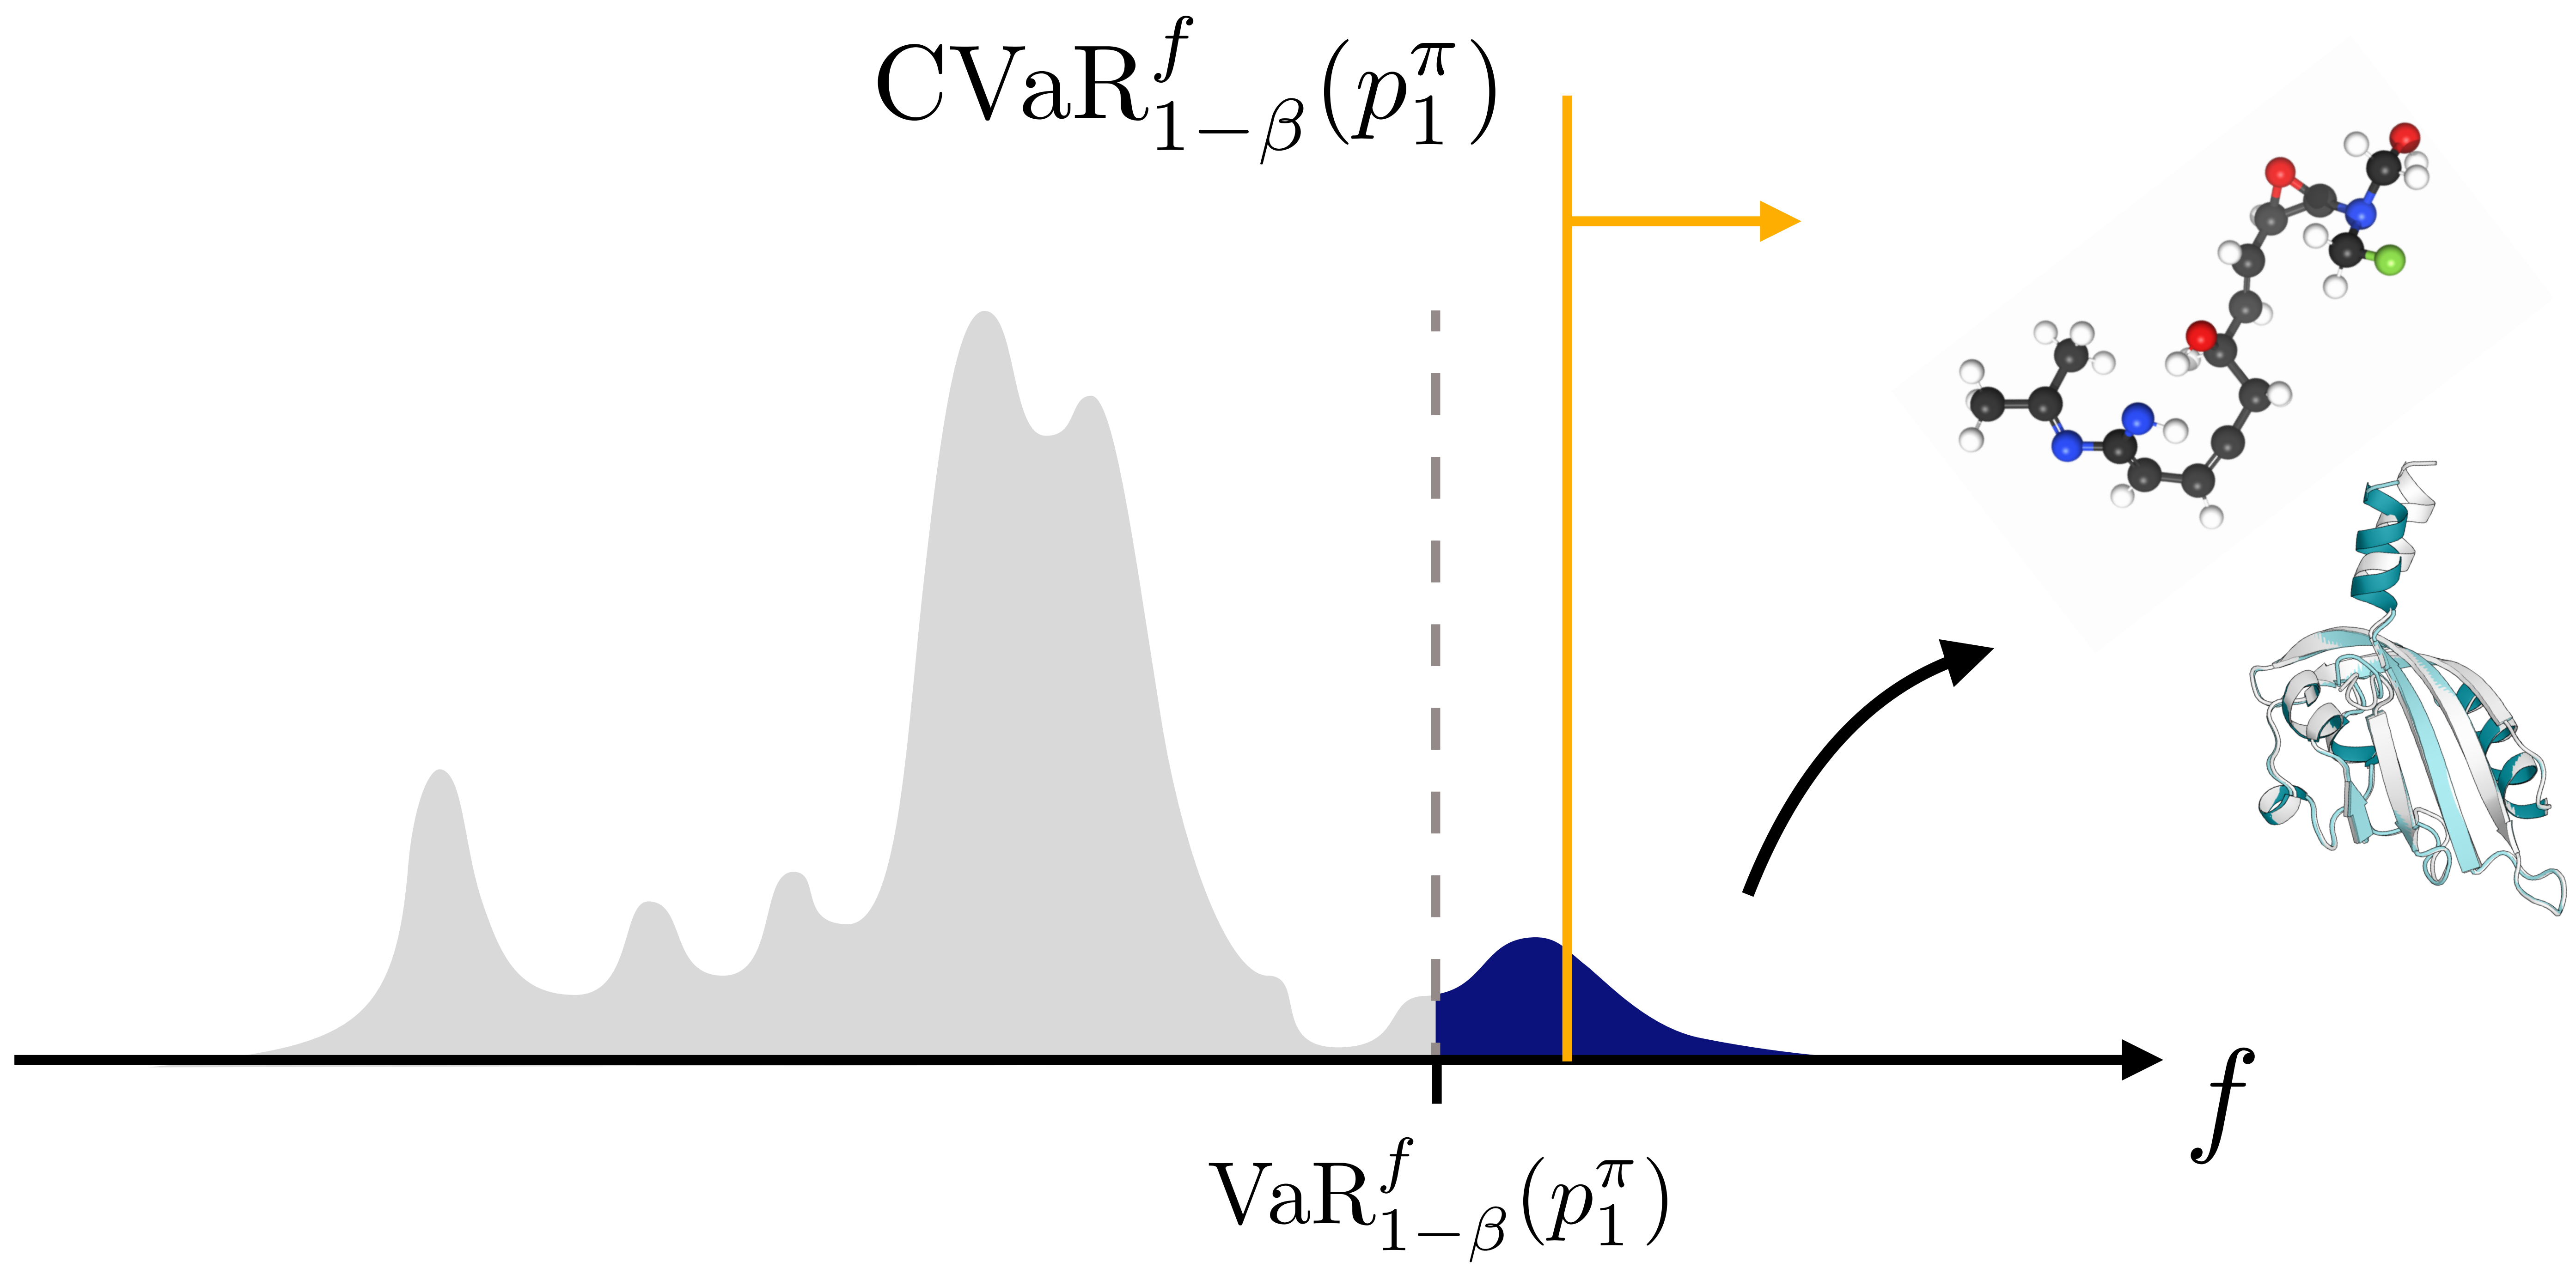

$$\mathrm{CVaR}^f_{1-\beta}(p^\pi_1) := \mathbb{E}_{x\sim p^\pi_1}[f(x) \mid f(x) \geq \mathrm{VaR}^f_{1-\beta}(p^\pi_1)]$$

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

Reward Tail-Aware Flow Steering for Discovery

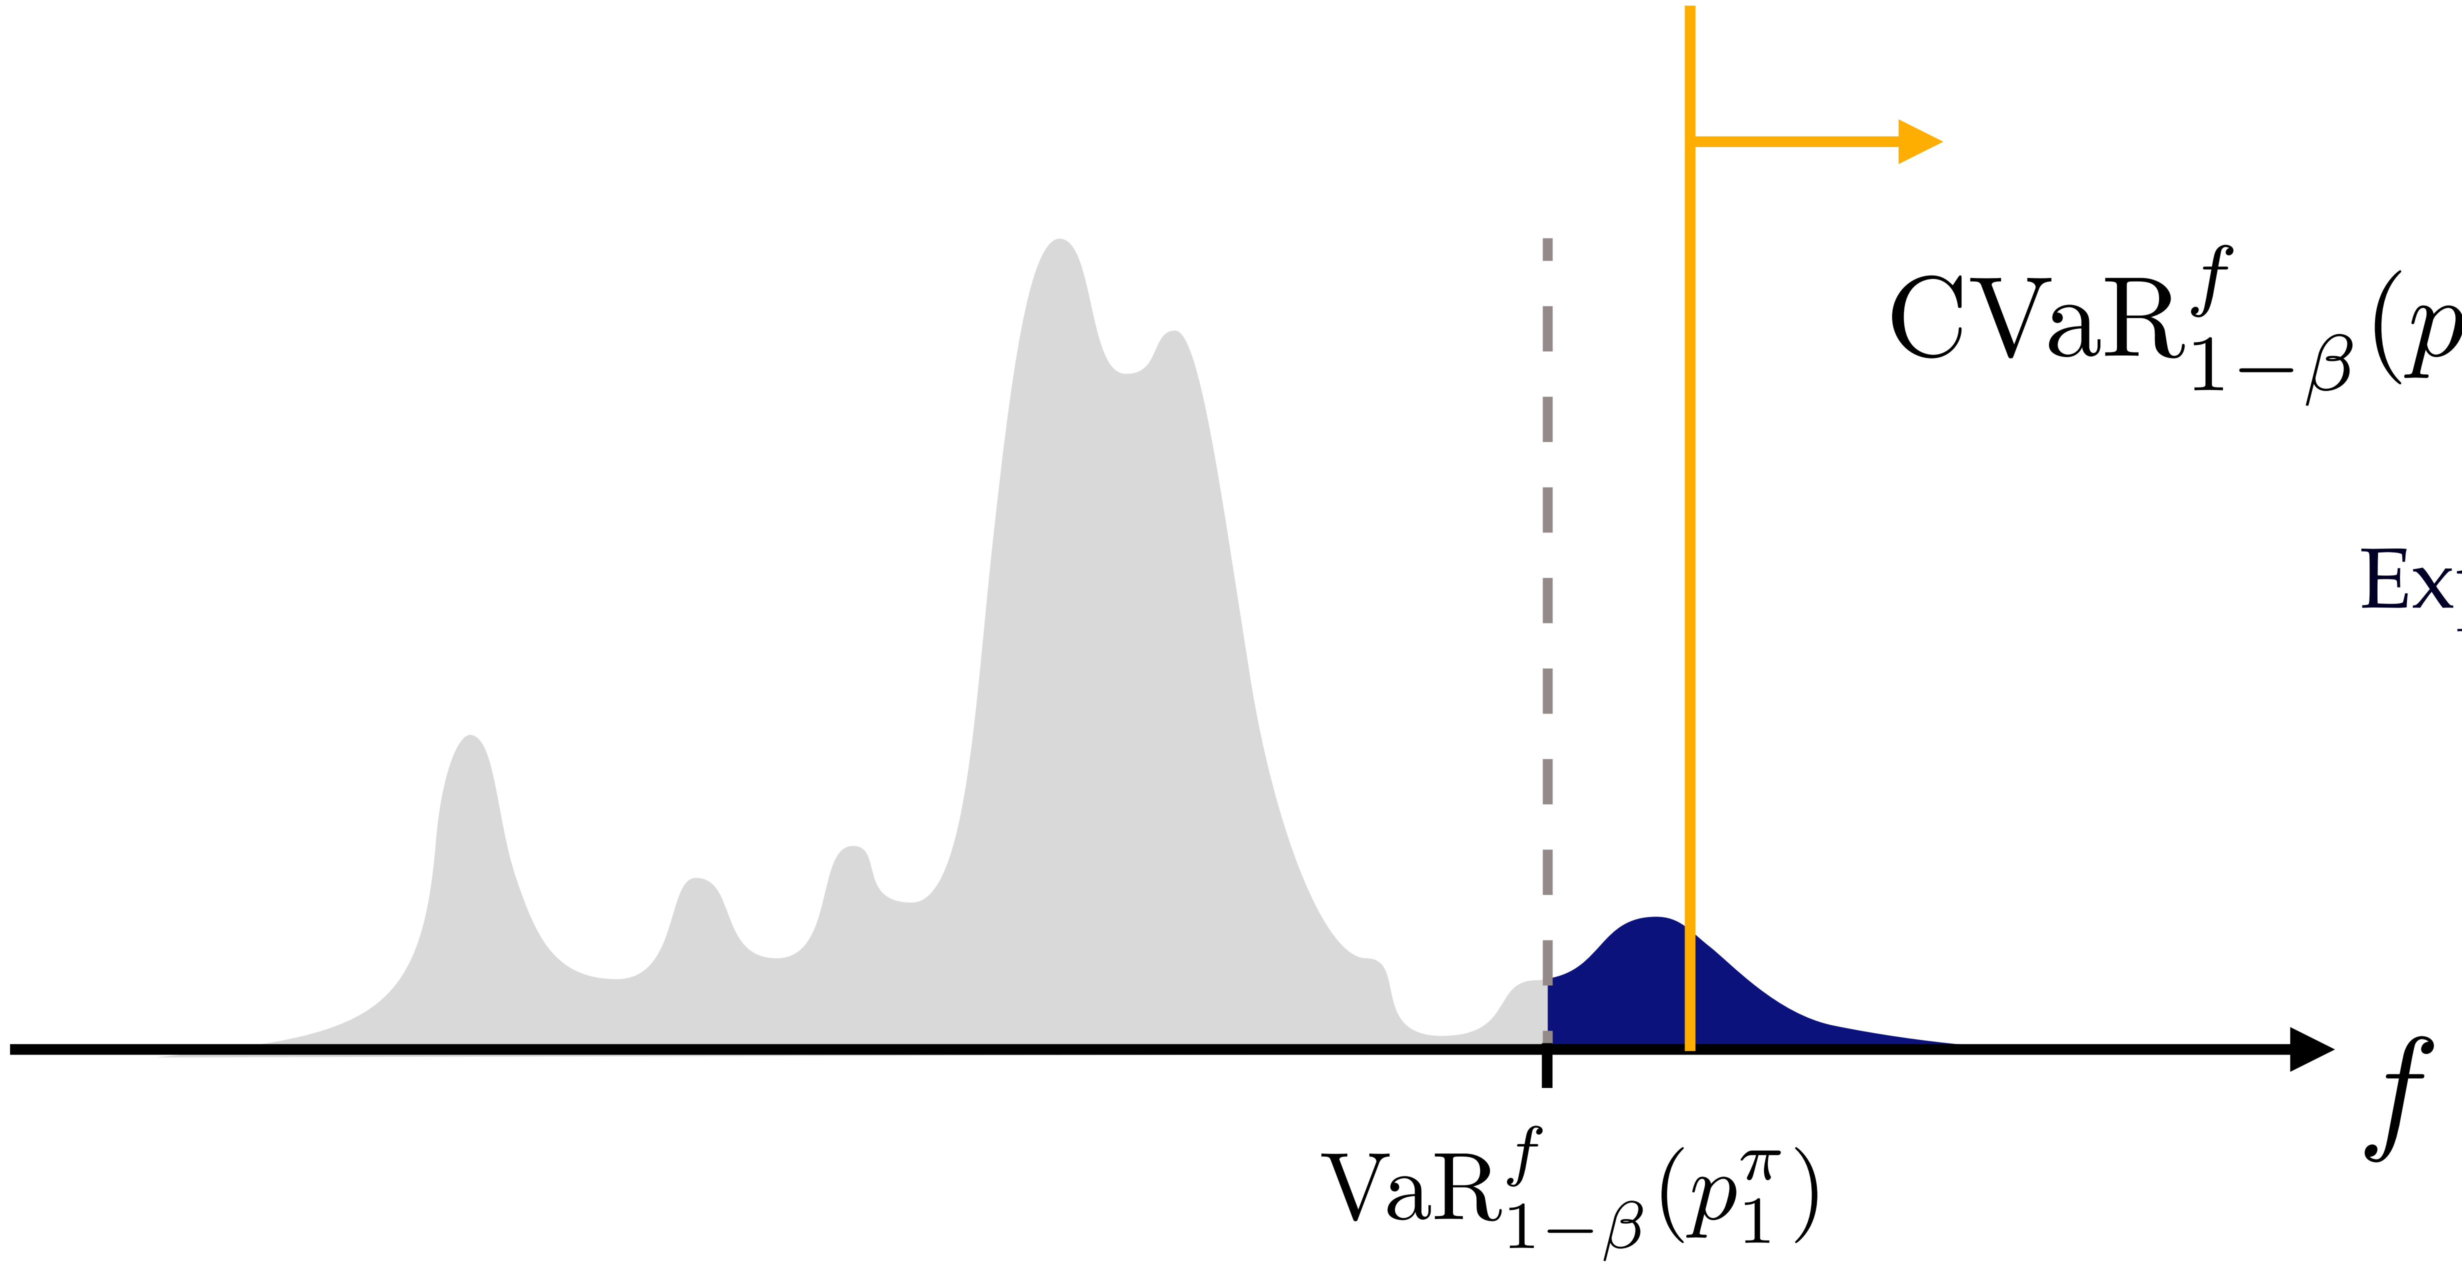

$$\mathrm{CVaR}^{f}_{1-\beta}(p^{\pi}_{1}) := \mathbb{E}_{x\sim p^{\pi}_{1}}[f(x) \mid f(x) \geq \mathrm{VaR}^{f}_{1-\beta}(p^{\pi}_{1})]$$

Expected value of top $1-\beta$ quantile

Reward Tail-Aware Flow Steering for Discovery

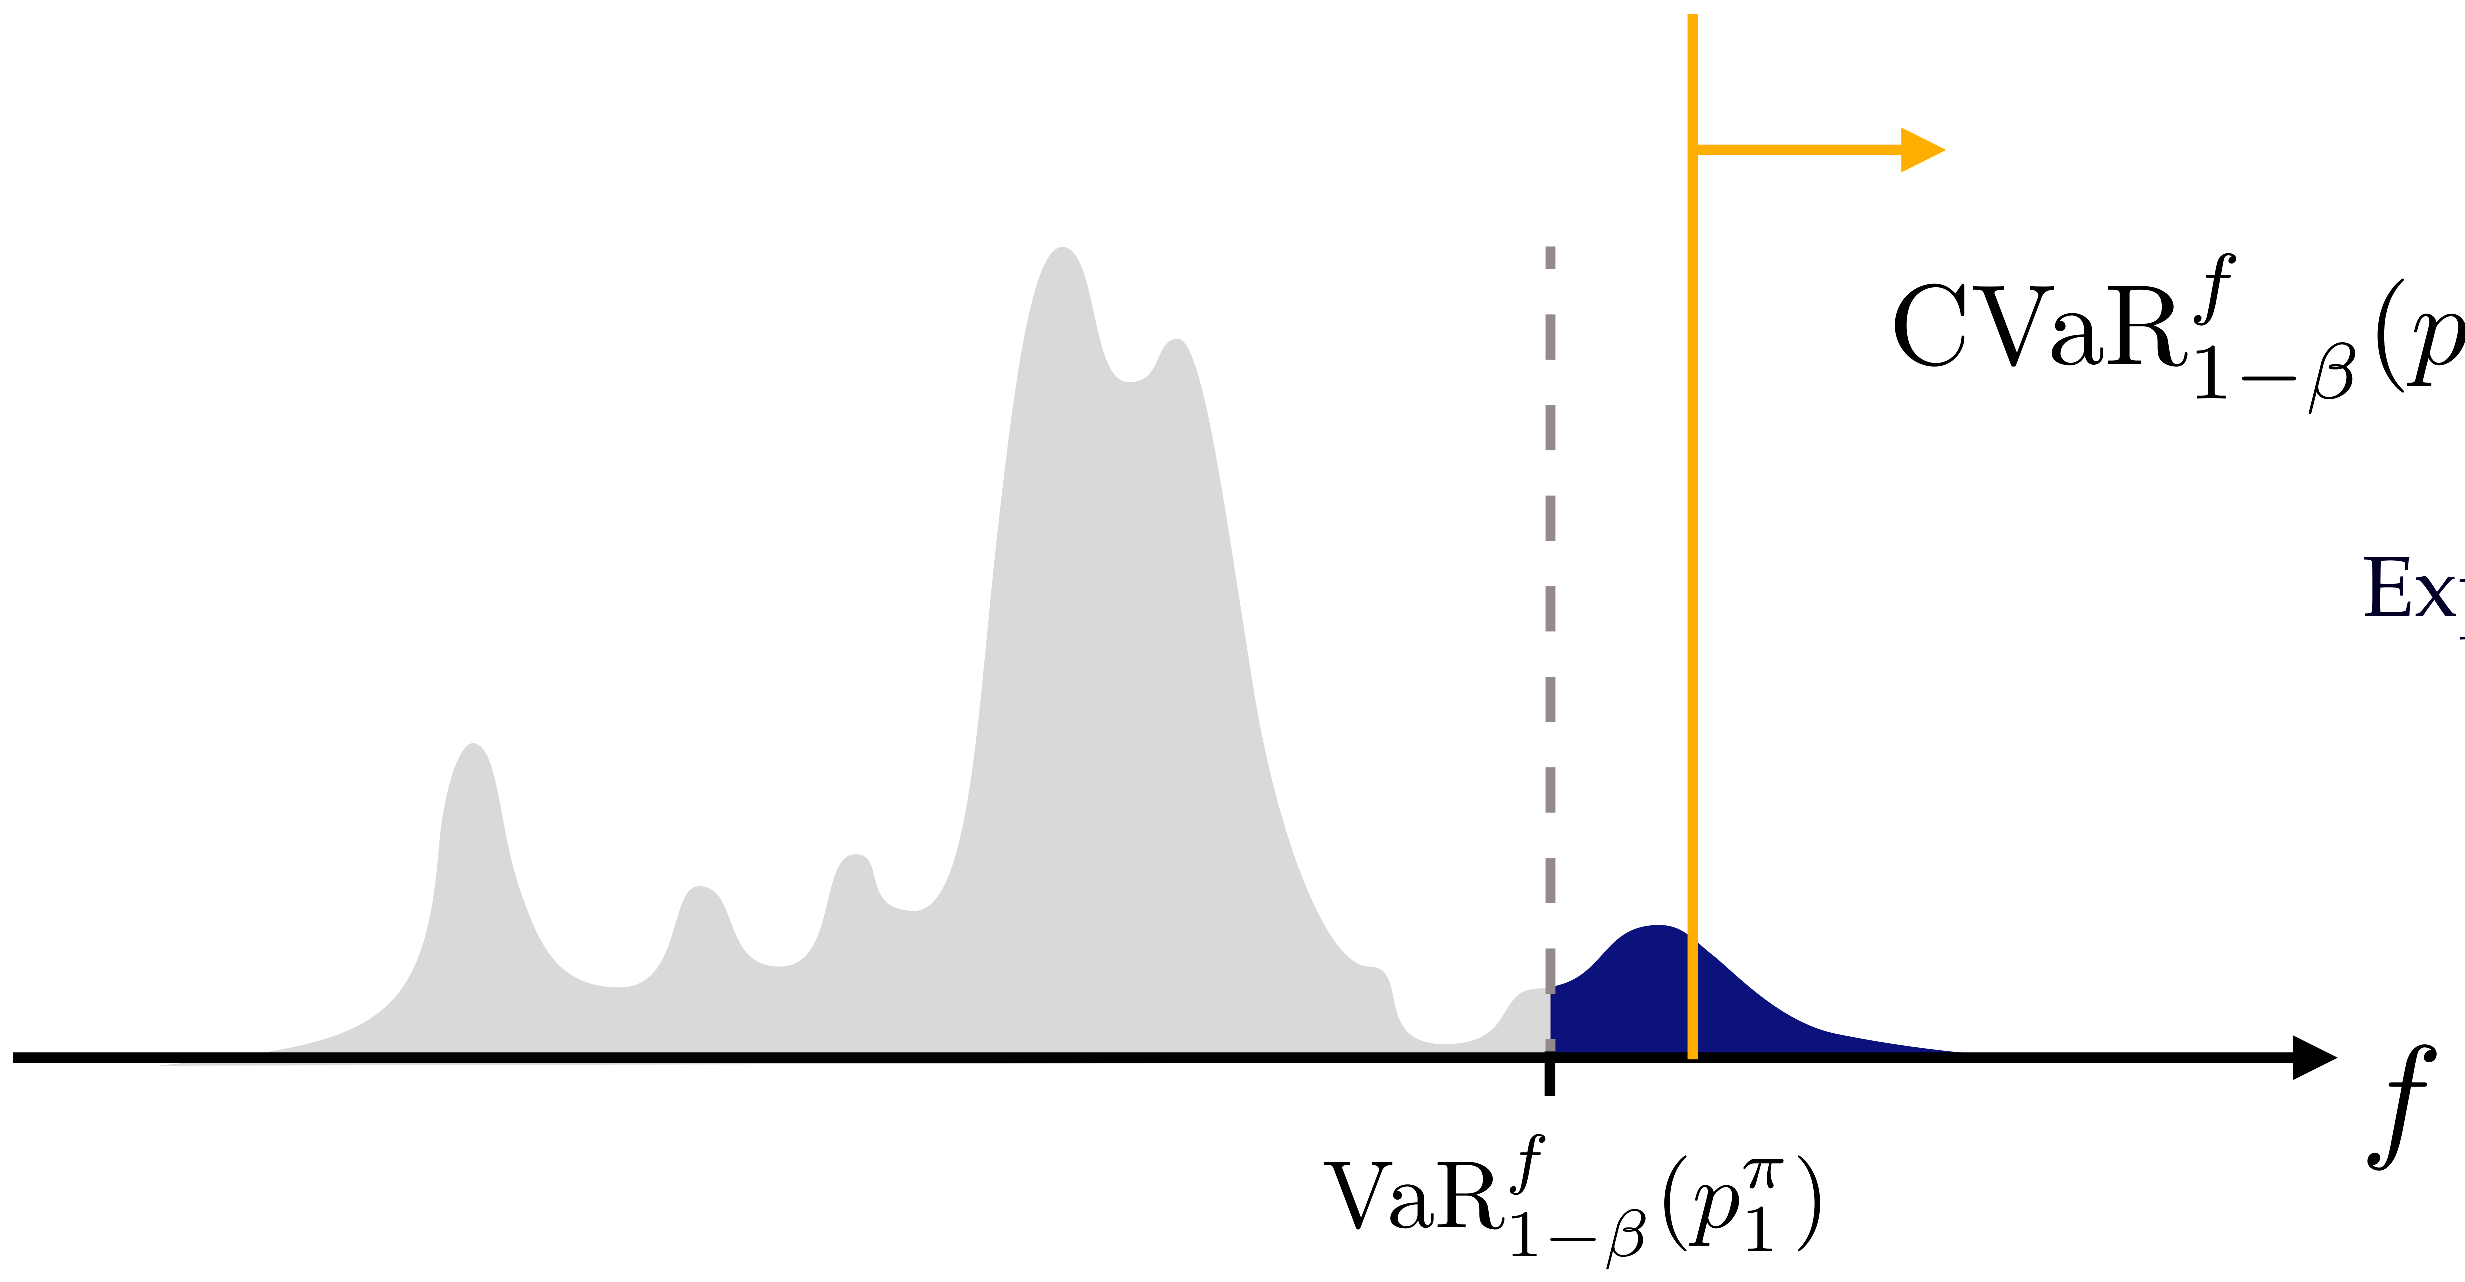

$$\mathrm{CVaR}^{f}_{1-\beta}(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \geq \mathrm{VaR}^{f}_{1-\beta}(p_1^\pi)]$$

Expected value of top $1-\beta$ quantile

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \arg\max_{\pi} \ \mathrm{CVaR}_{1-\beta}(p_1^\pi)$$

Reward Tail-Aware Flow Steering for Discovery

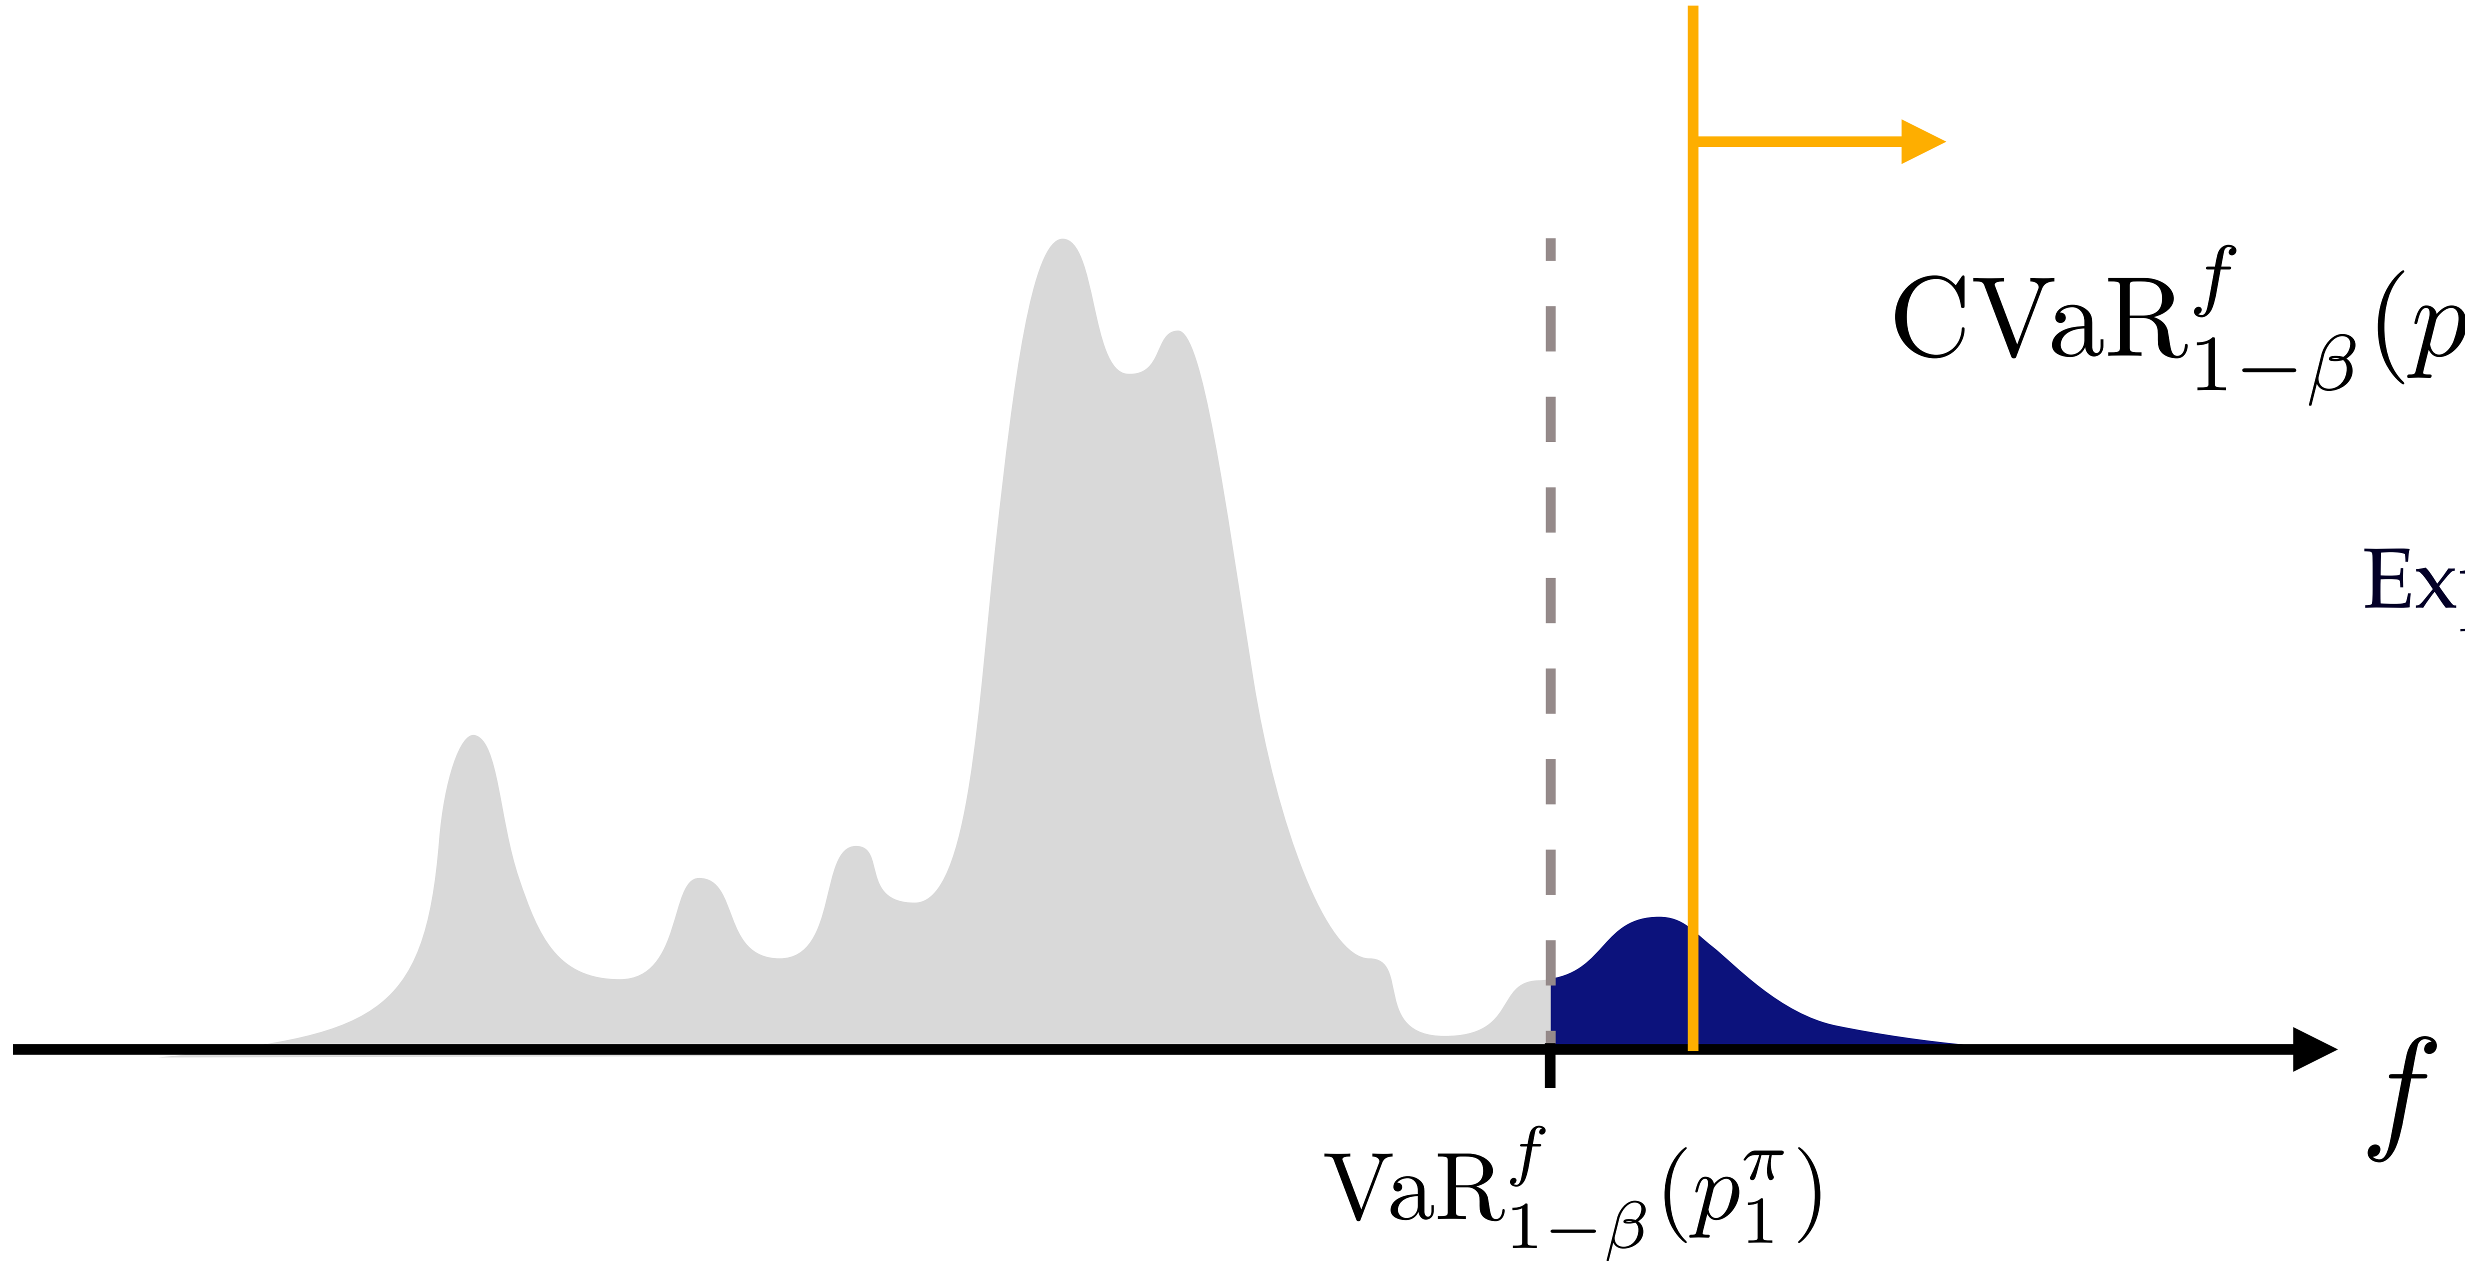

$$\mathrm{CVaR}^f_{1-\beta}(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \geq \mathrm{VaR}^f_{1-\beta}(p_1^\pi)]$$

Expected value of top $1-\beta$ quantile

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \operatorname*{argmax}_{\pi} \ \mathrm{CVaR}_{1-\beta}(p_1^\pi)$$

Sacrifices average sample reward for top samples

Reward Tail-Aware Flow Steering for Discovery

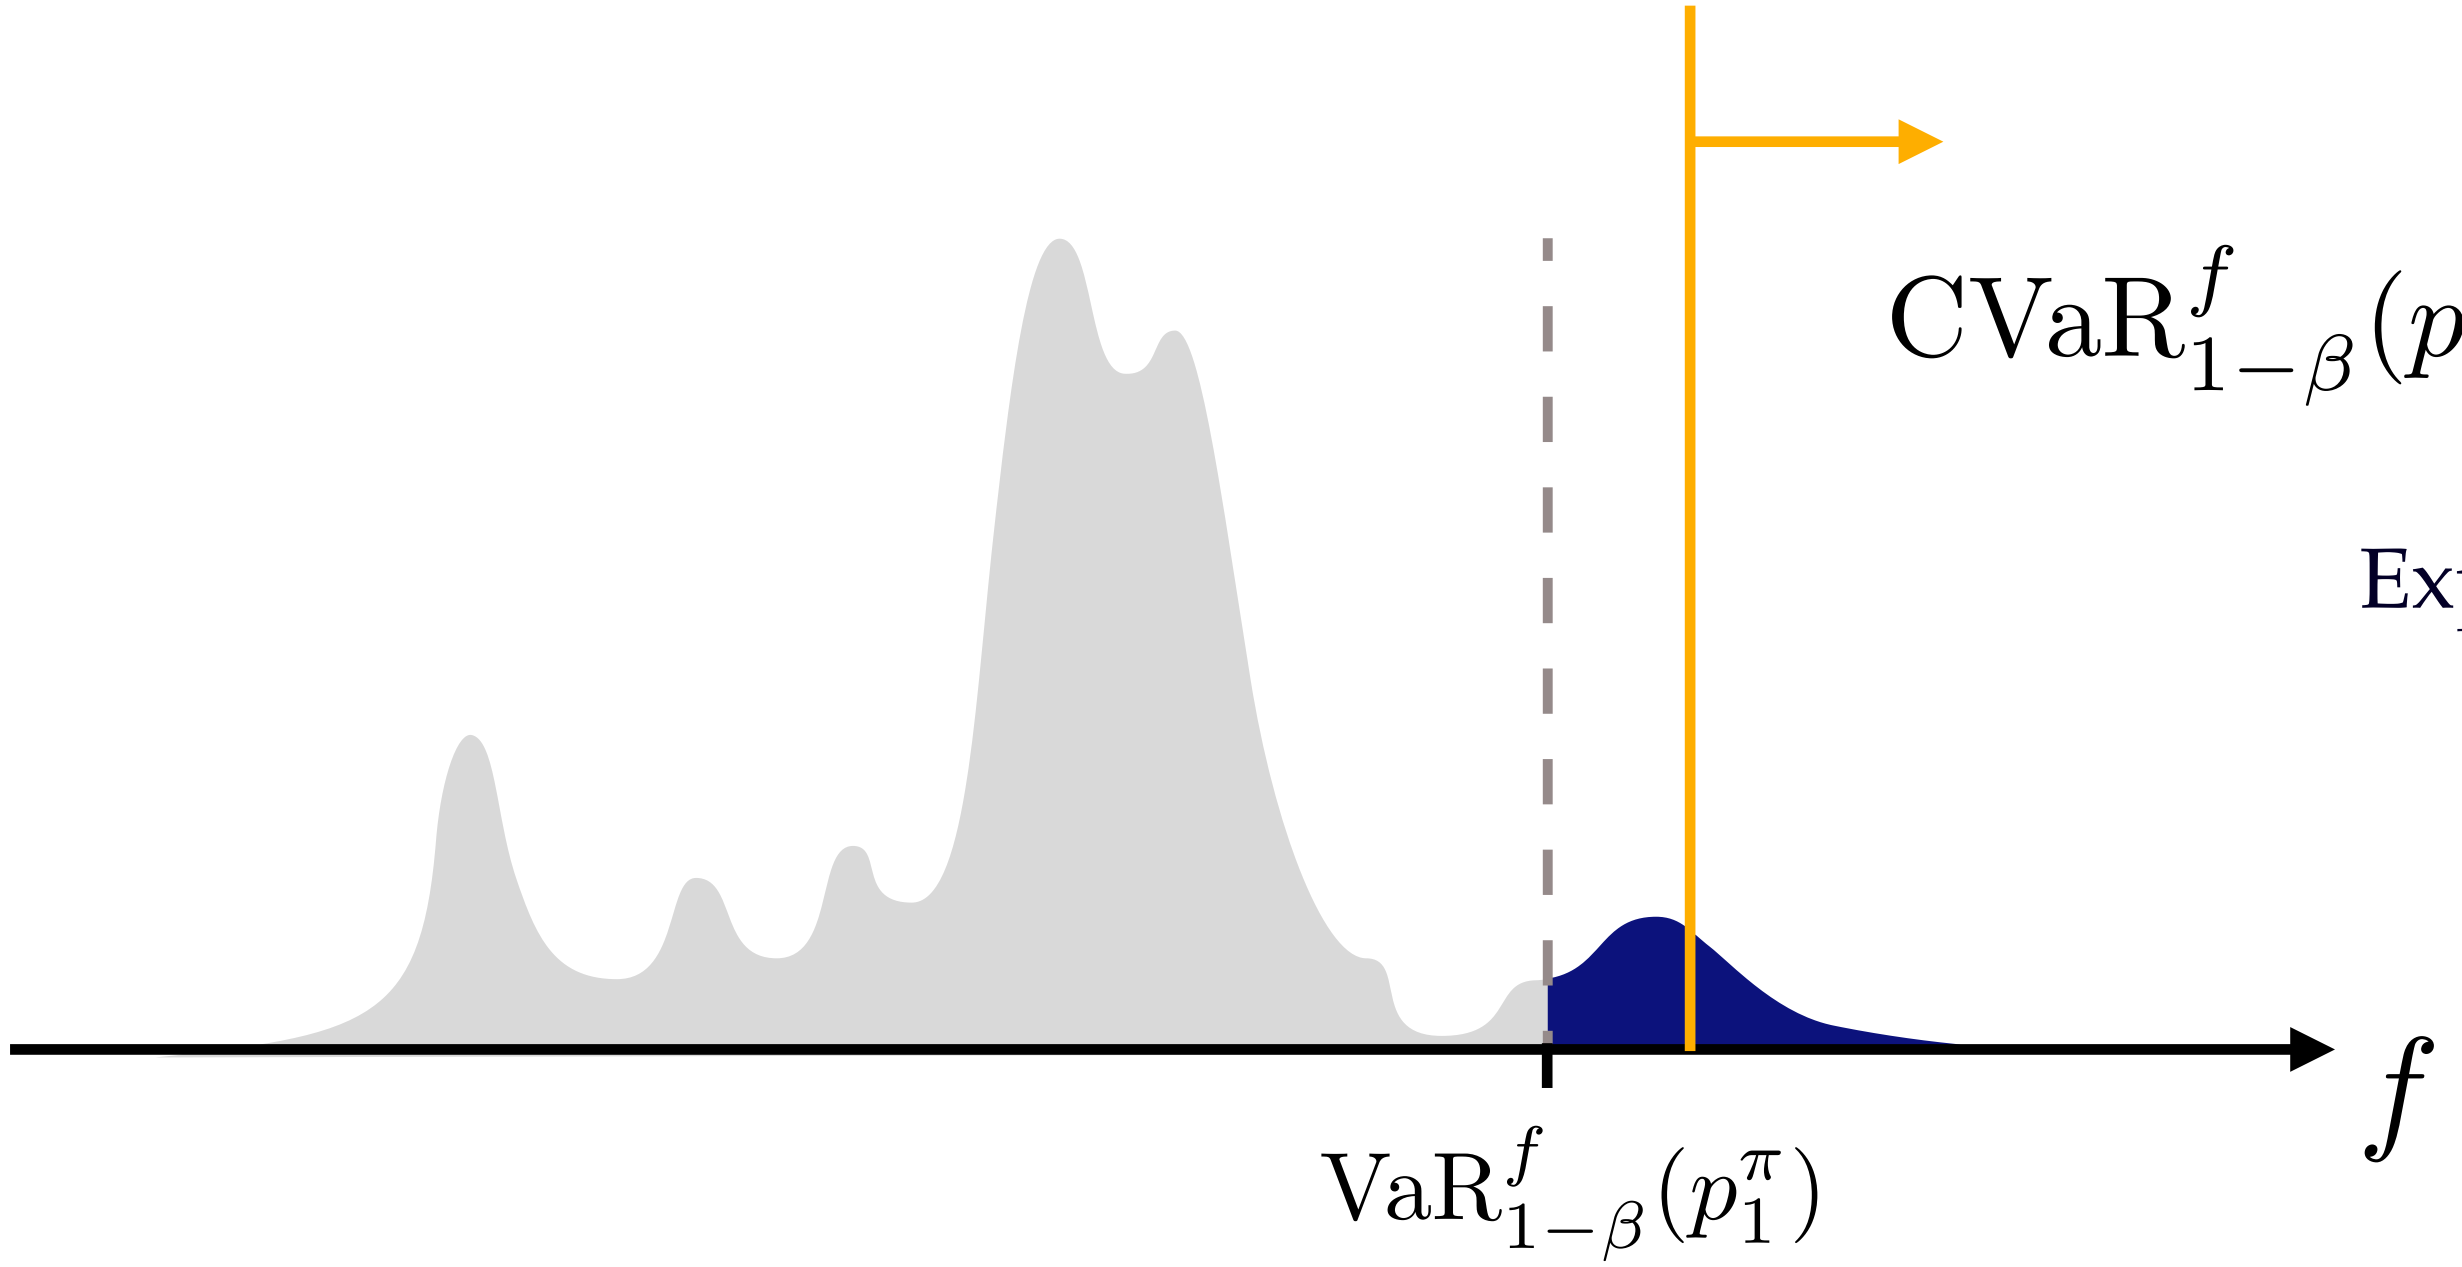

$$\mathrm{CVaR}^{f}_{1-\beta}(p^{\pi}_1) := \mathbb{E}_{x\sim p^{\pi}_1}[f(x) \mid f(x) \geq \mathrm{VaR}^{f}_{1-\beta}(p^{\pi}_1)]$$

Expected value of top $1-\beta$ quantile

Tail-Aware Reward-Guided Flow Fine-Tuning

$$\pi^* \in \underset{\pi}{\mathrm{argmax}} \; \mathrm{CVaR}_{1-\beta}(p^{\pi}_1)$$

Sacrifices average sample reward for top samples

Higher chances to discover truly exceptional designs

Reward Tail-Aware Flow Steering for Discovery

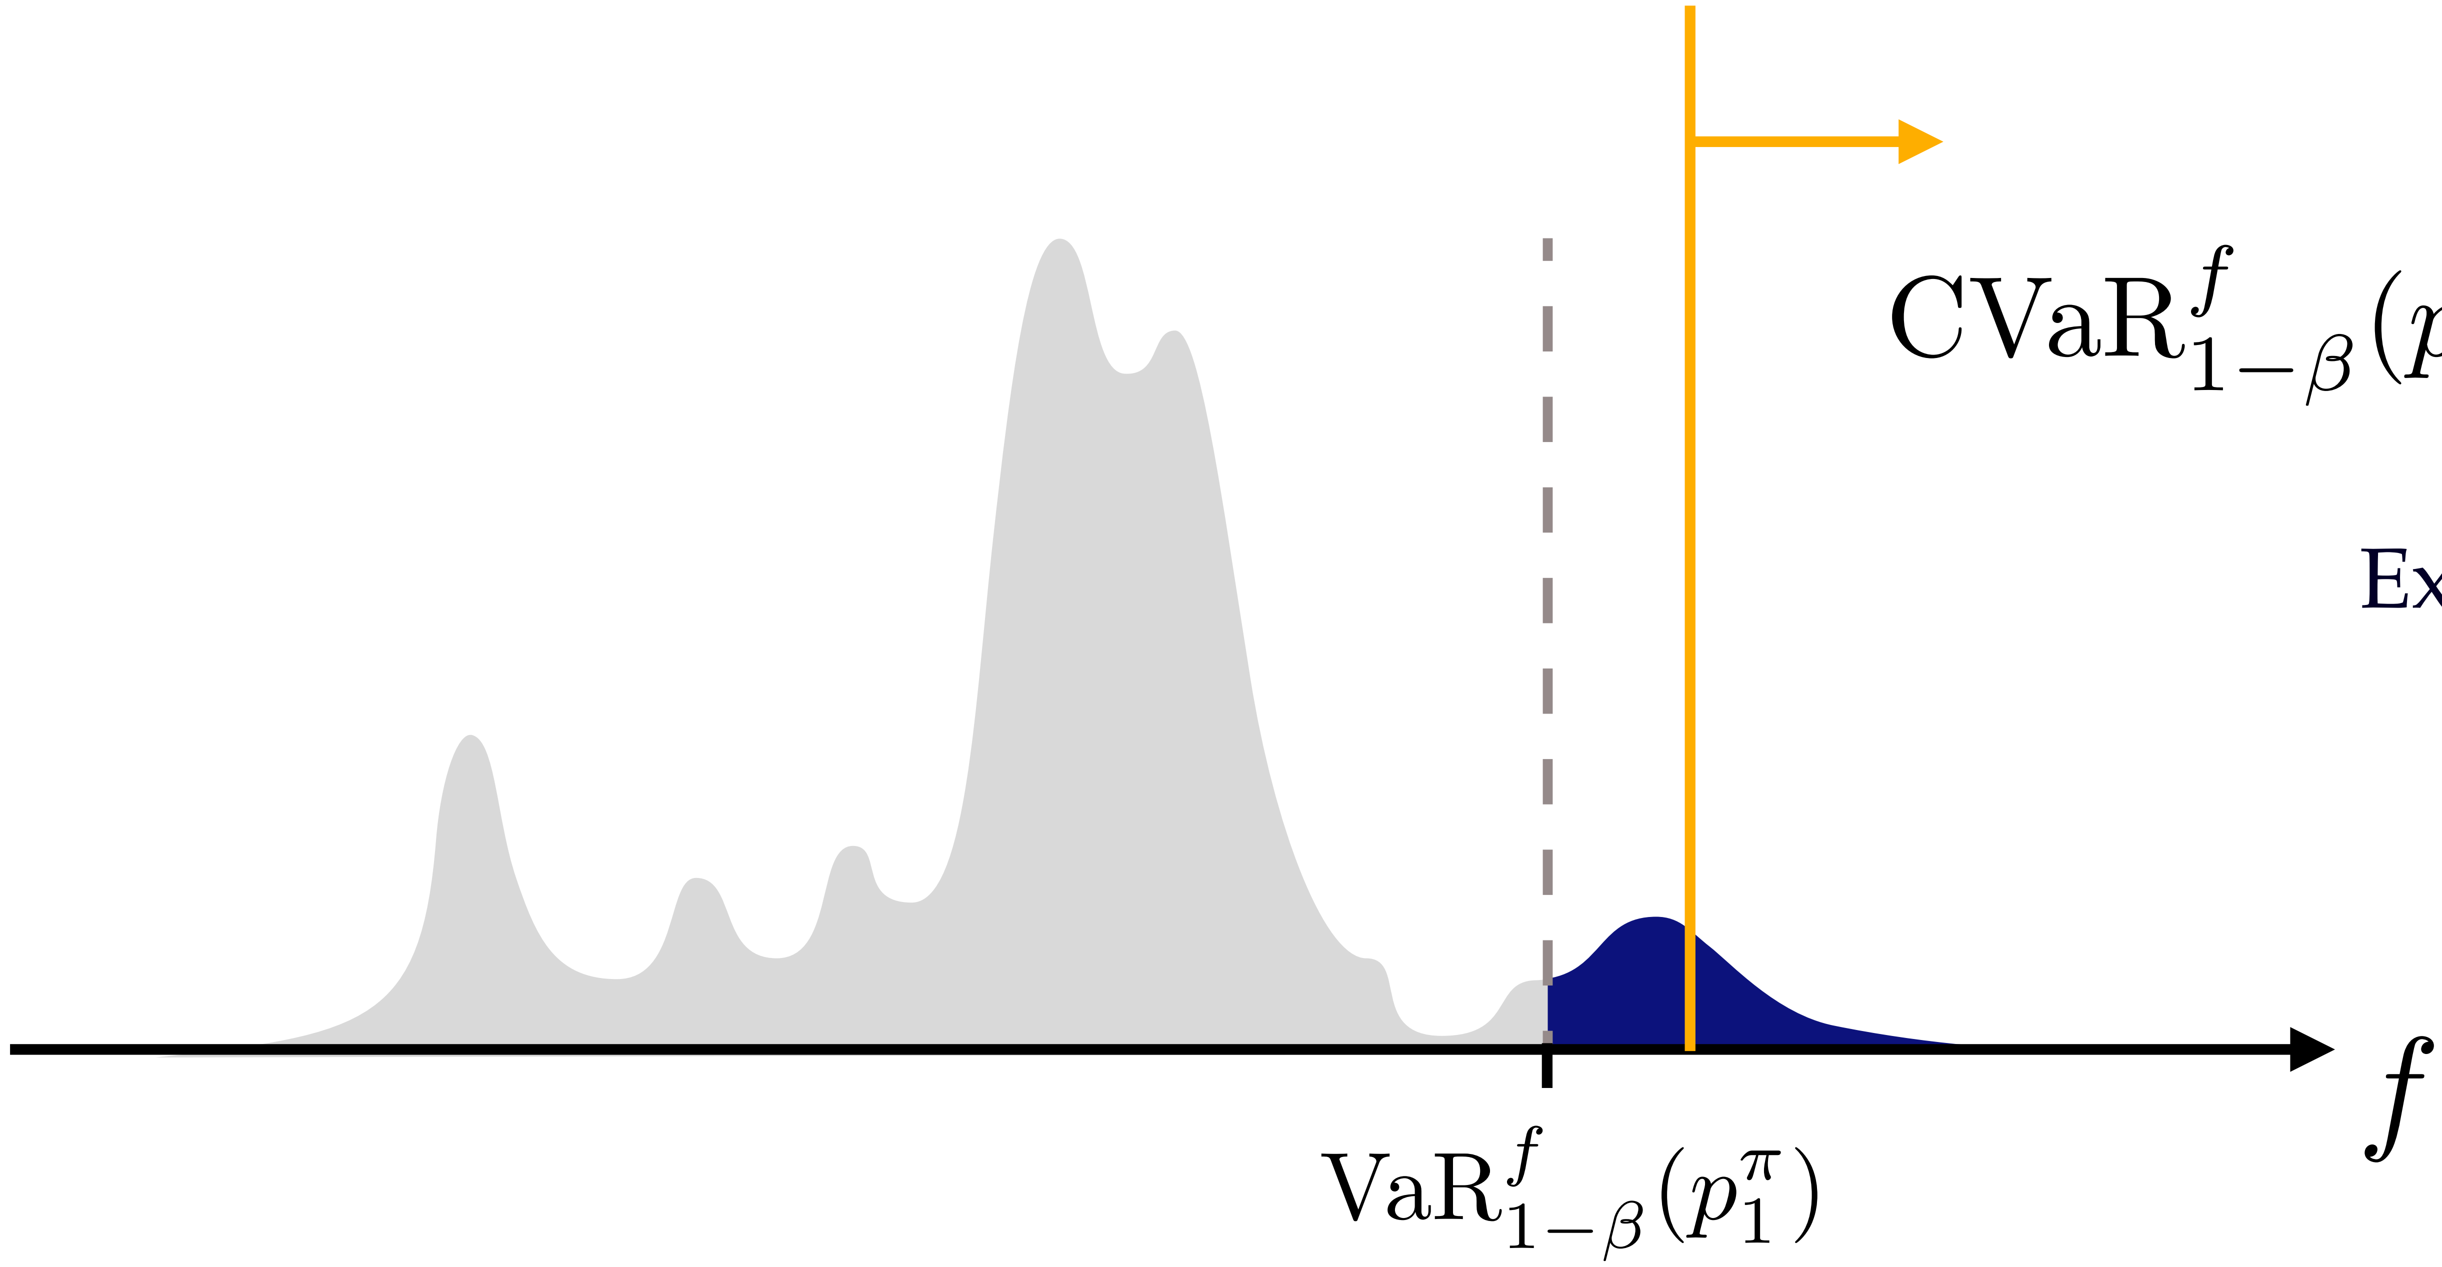

$$\mathrm{CVaR}^f_{1-\beta}(p_1^\pi) := \mathbb{E}_{x\sim p_1^\pi}[f(x) \mid f(x) \geq \mathrm{VaR}^f_{1-\beta}(p_1^\pi)]$$

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**

↓ **Sacrifices average sample reward for top samples**

↓ **Higher chances to discover truly exceptional designs**

Reward Tail-Aware Flow Steering as Distributional Fine-Tuning

Expect Reward

$$\mathbb{E}_{x \sim p_1^\pi} [f(x)]$$

Linear functional

→

Conditional Value at Risk

$$\mathrm{CVaR}^f_{1-\beta}(p_1^\pi)$$

Non-Linear functional

How to fine-tune a generative model to a **nonlinear** functional?

How to Maximize CVaR via RL

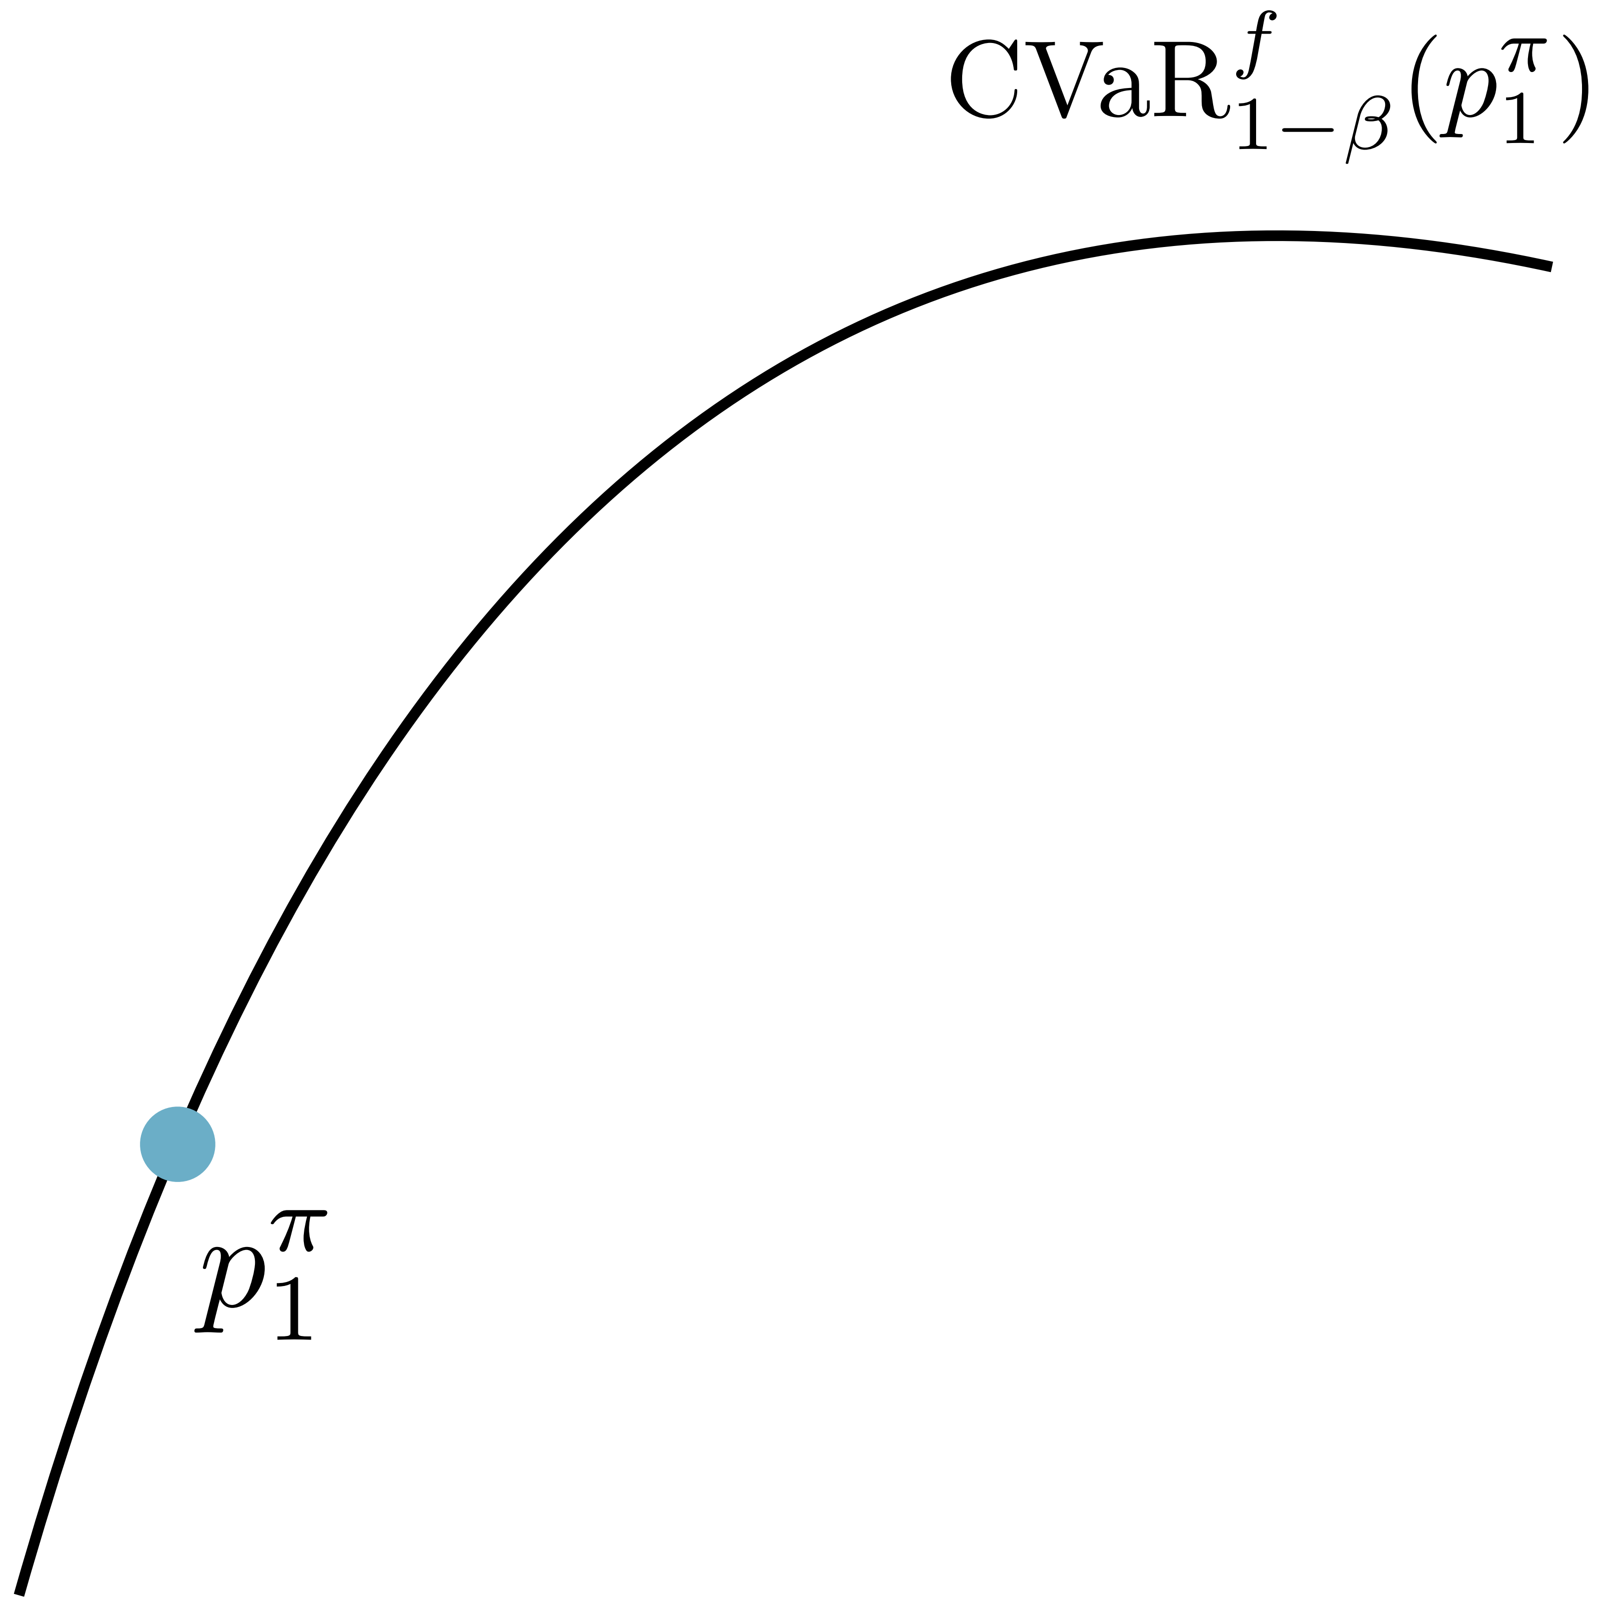

How to Maximize CVaR via RL

$$r := \frac{\delta \mathrm{CVaR}^{f}_{1-\beta}(p_1^\pi)}{\delta p_1^\pi}$$

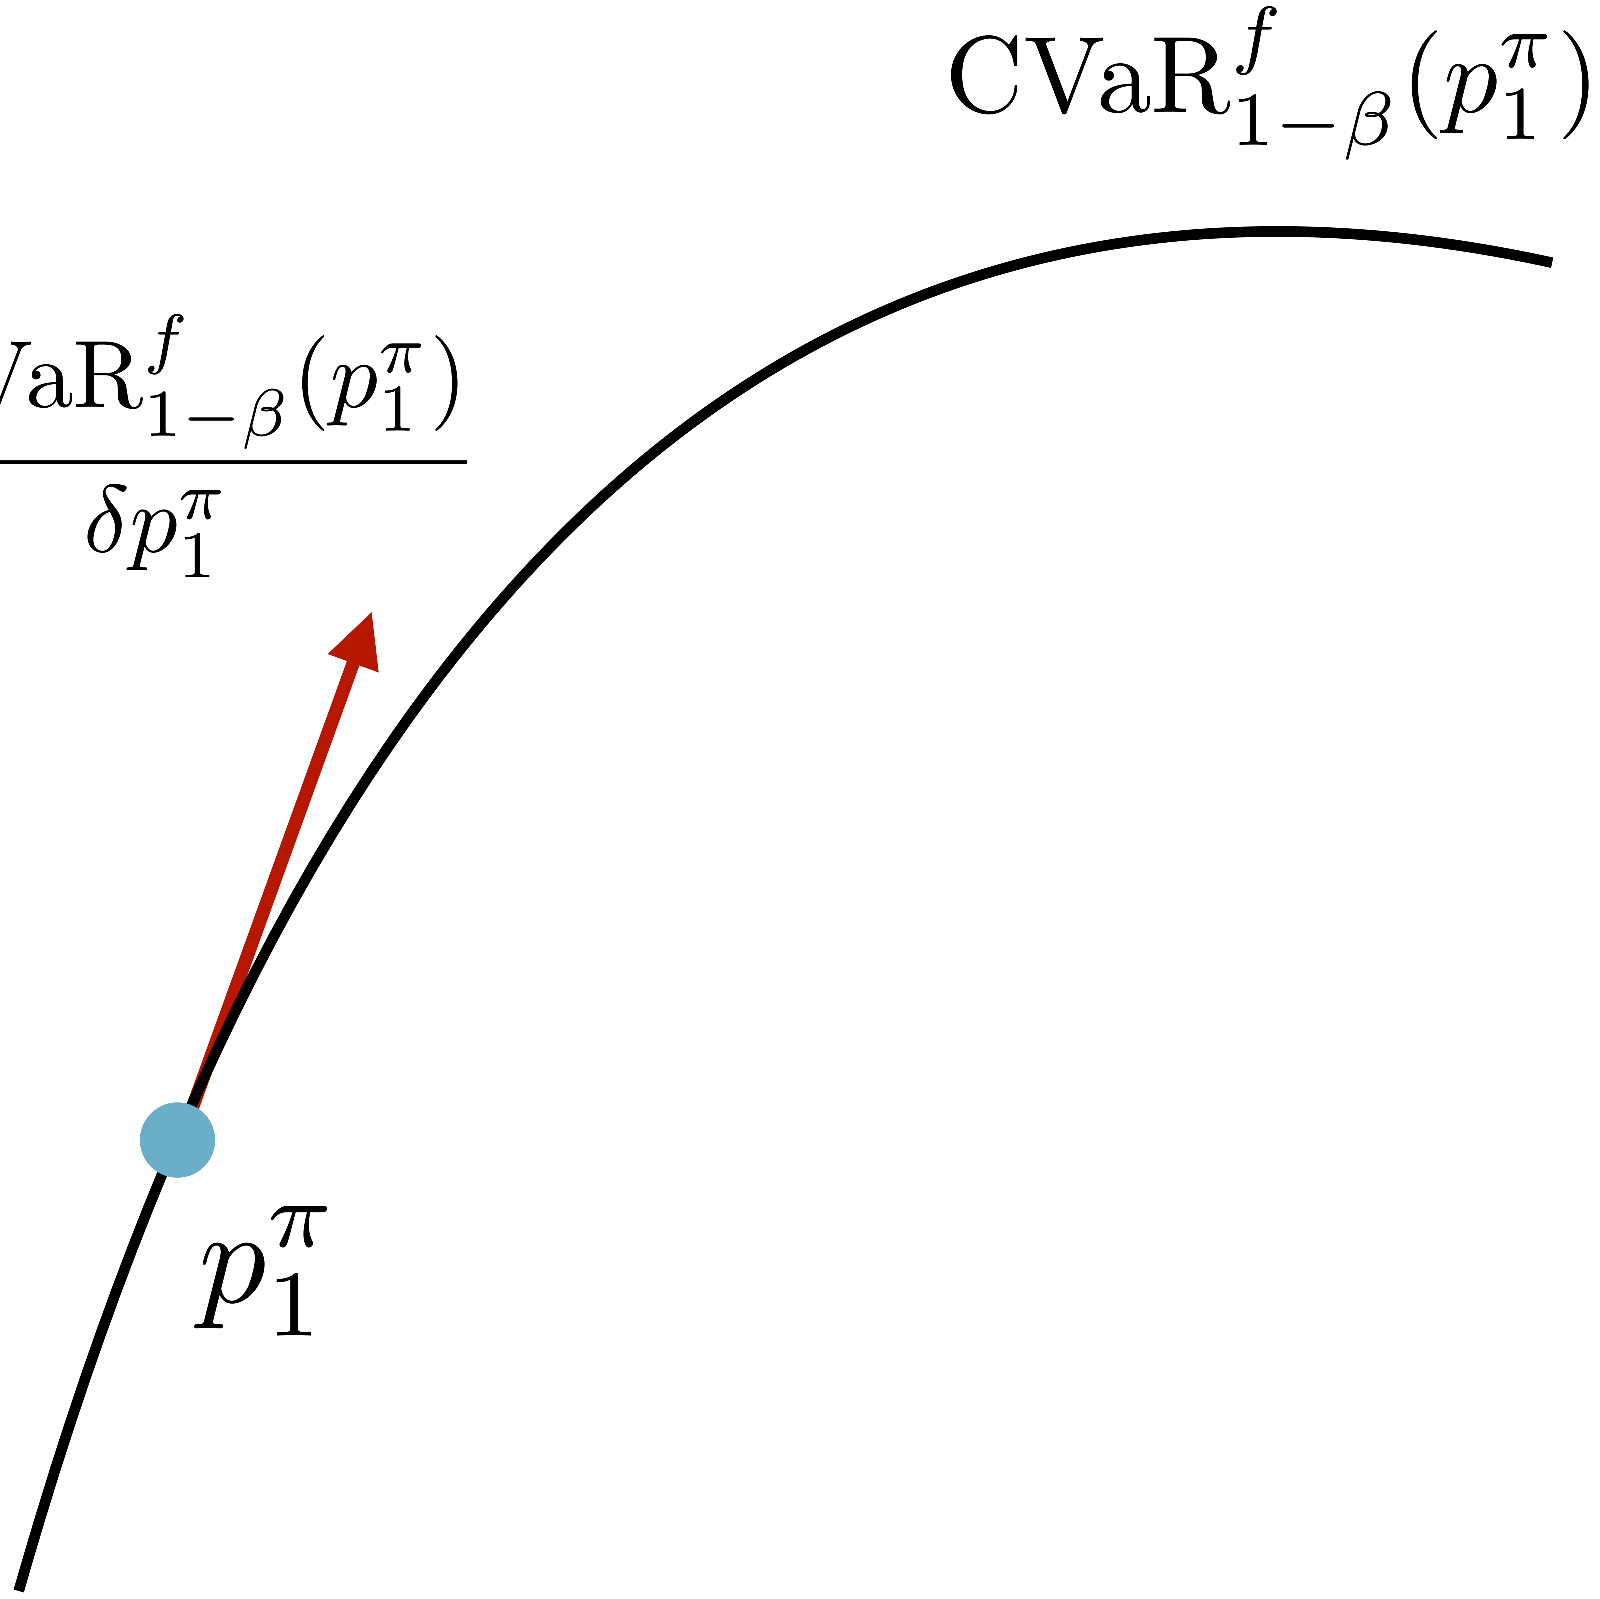

How to Maximize CVaR via RL

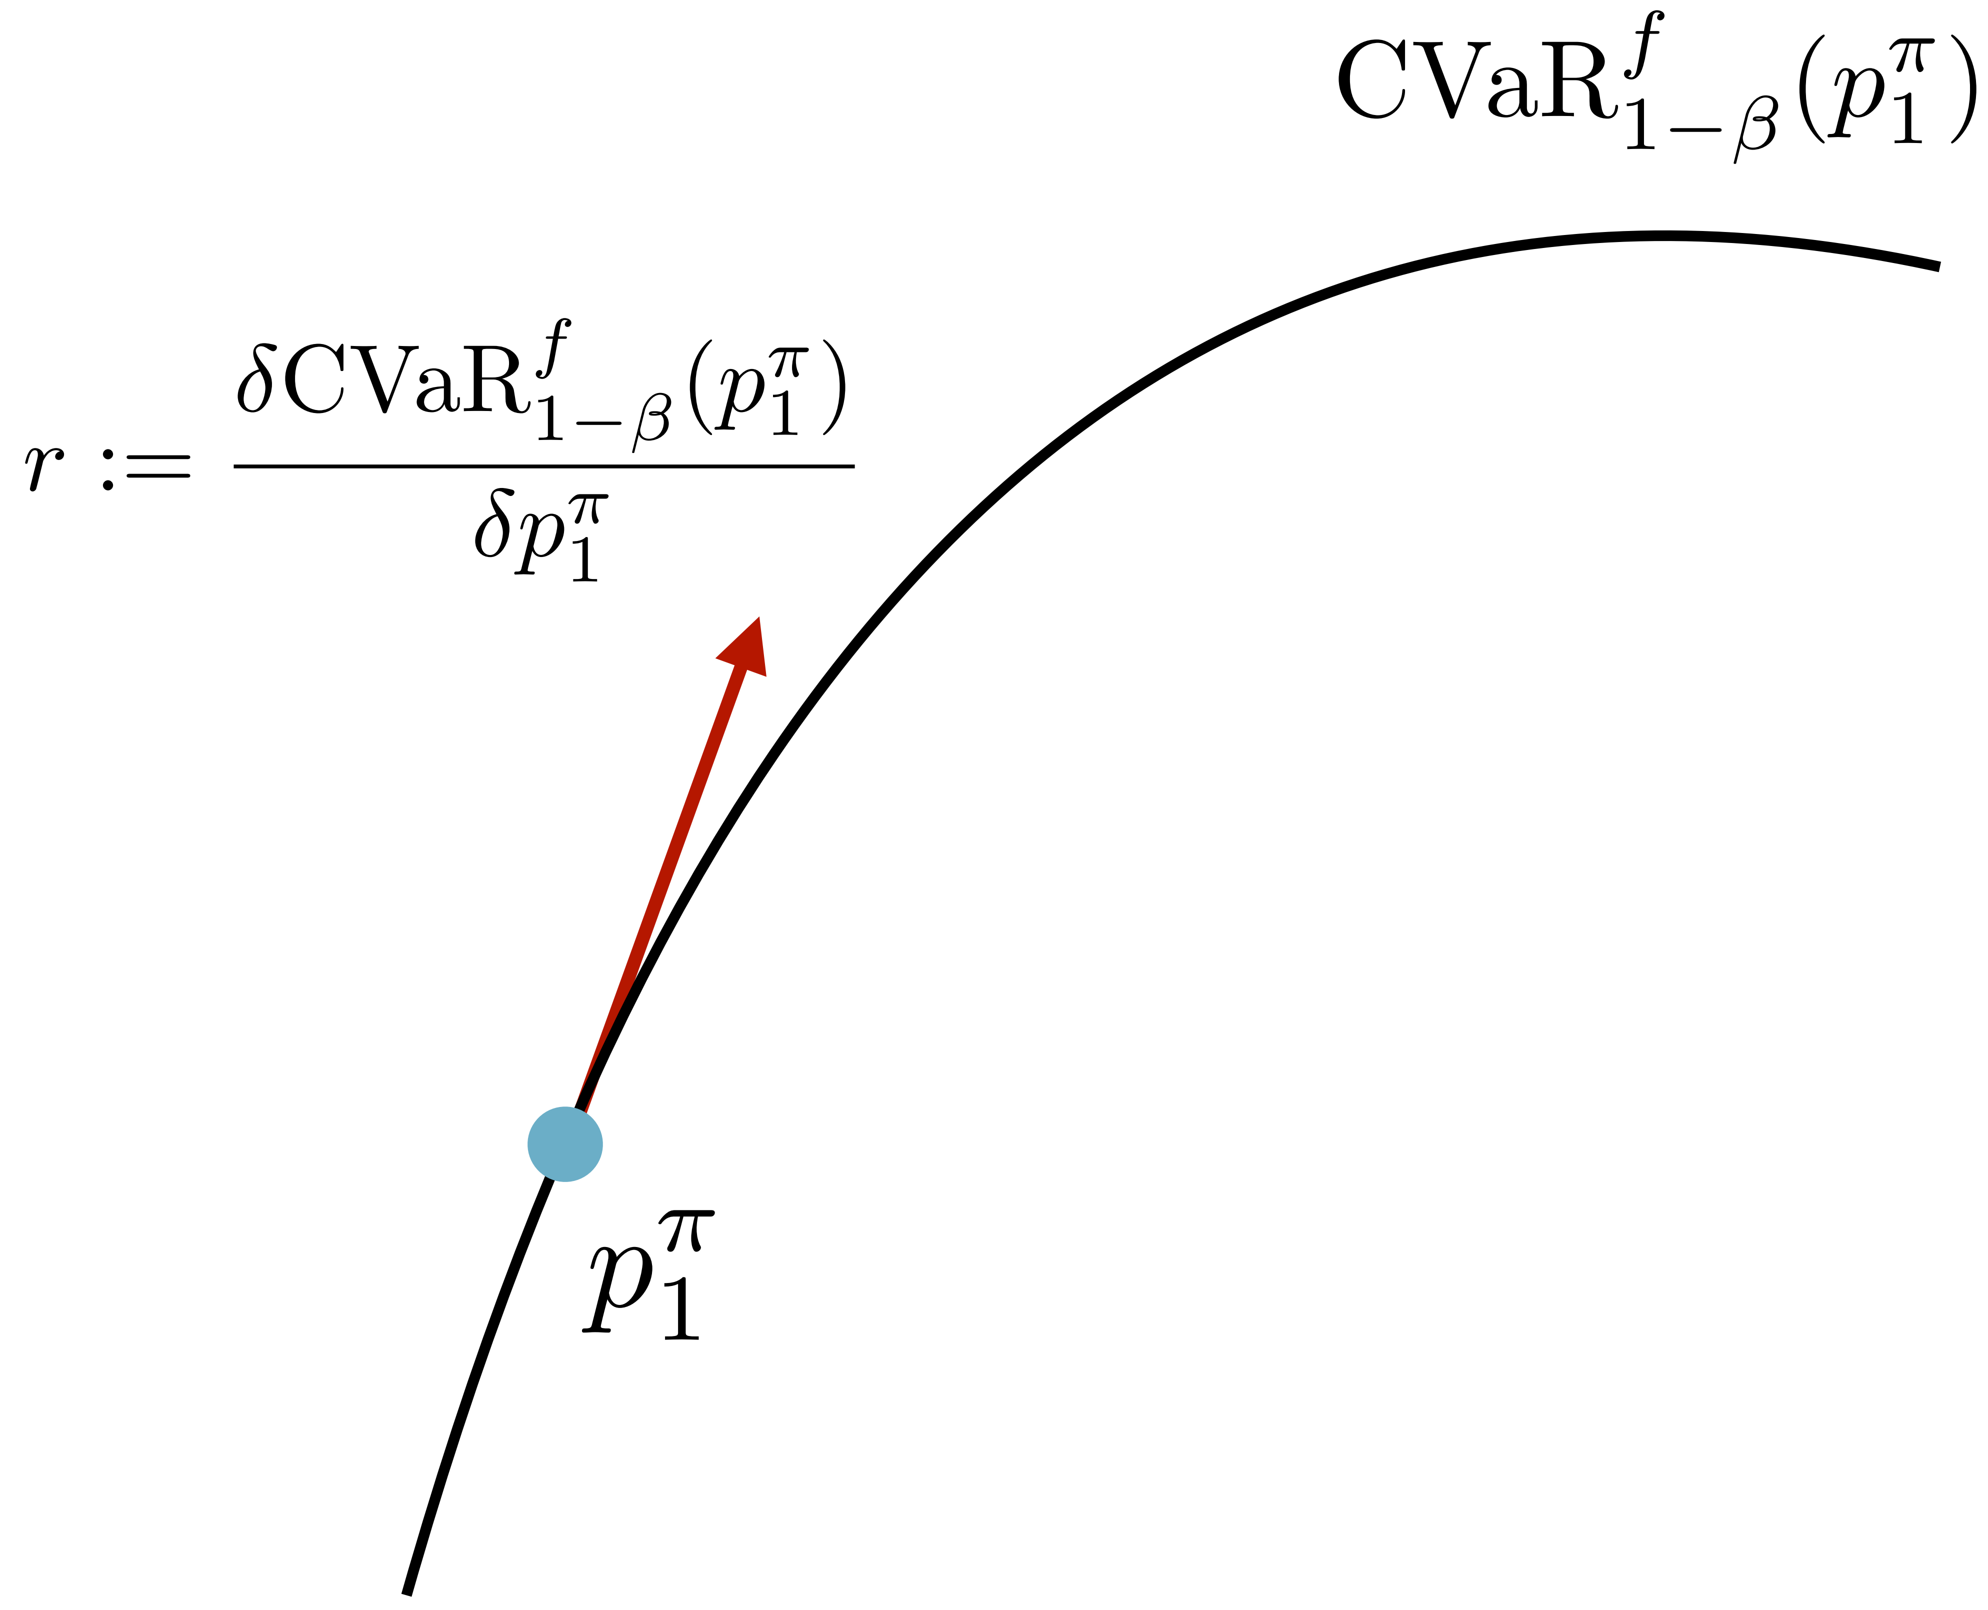

CVaR First Variation

$$r := \frac{\delta \mathrm{CVaR}^f_{1-\beta}(p_1^\pi)}{\delta p_1^\pi} = \frac{(f(x) - VaR^f_{1-\beta}(p_1^\pi))}{\beta}$$

↓

Can use RL fine-tuning

How to Maximize CVaR via RL

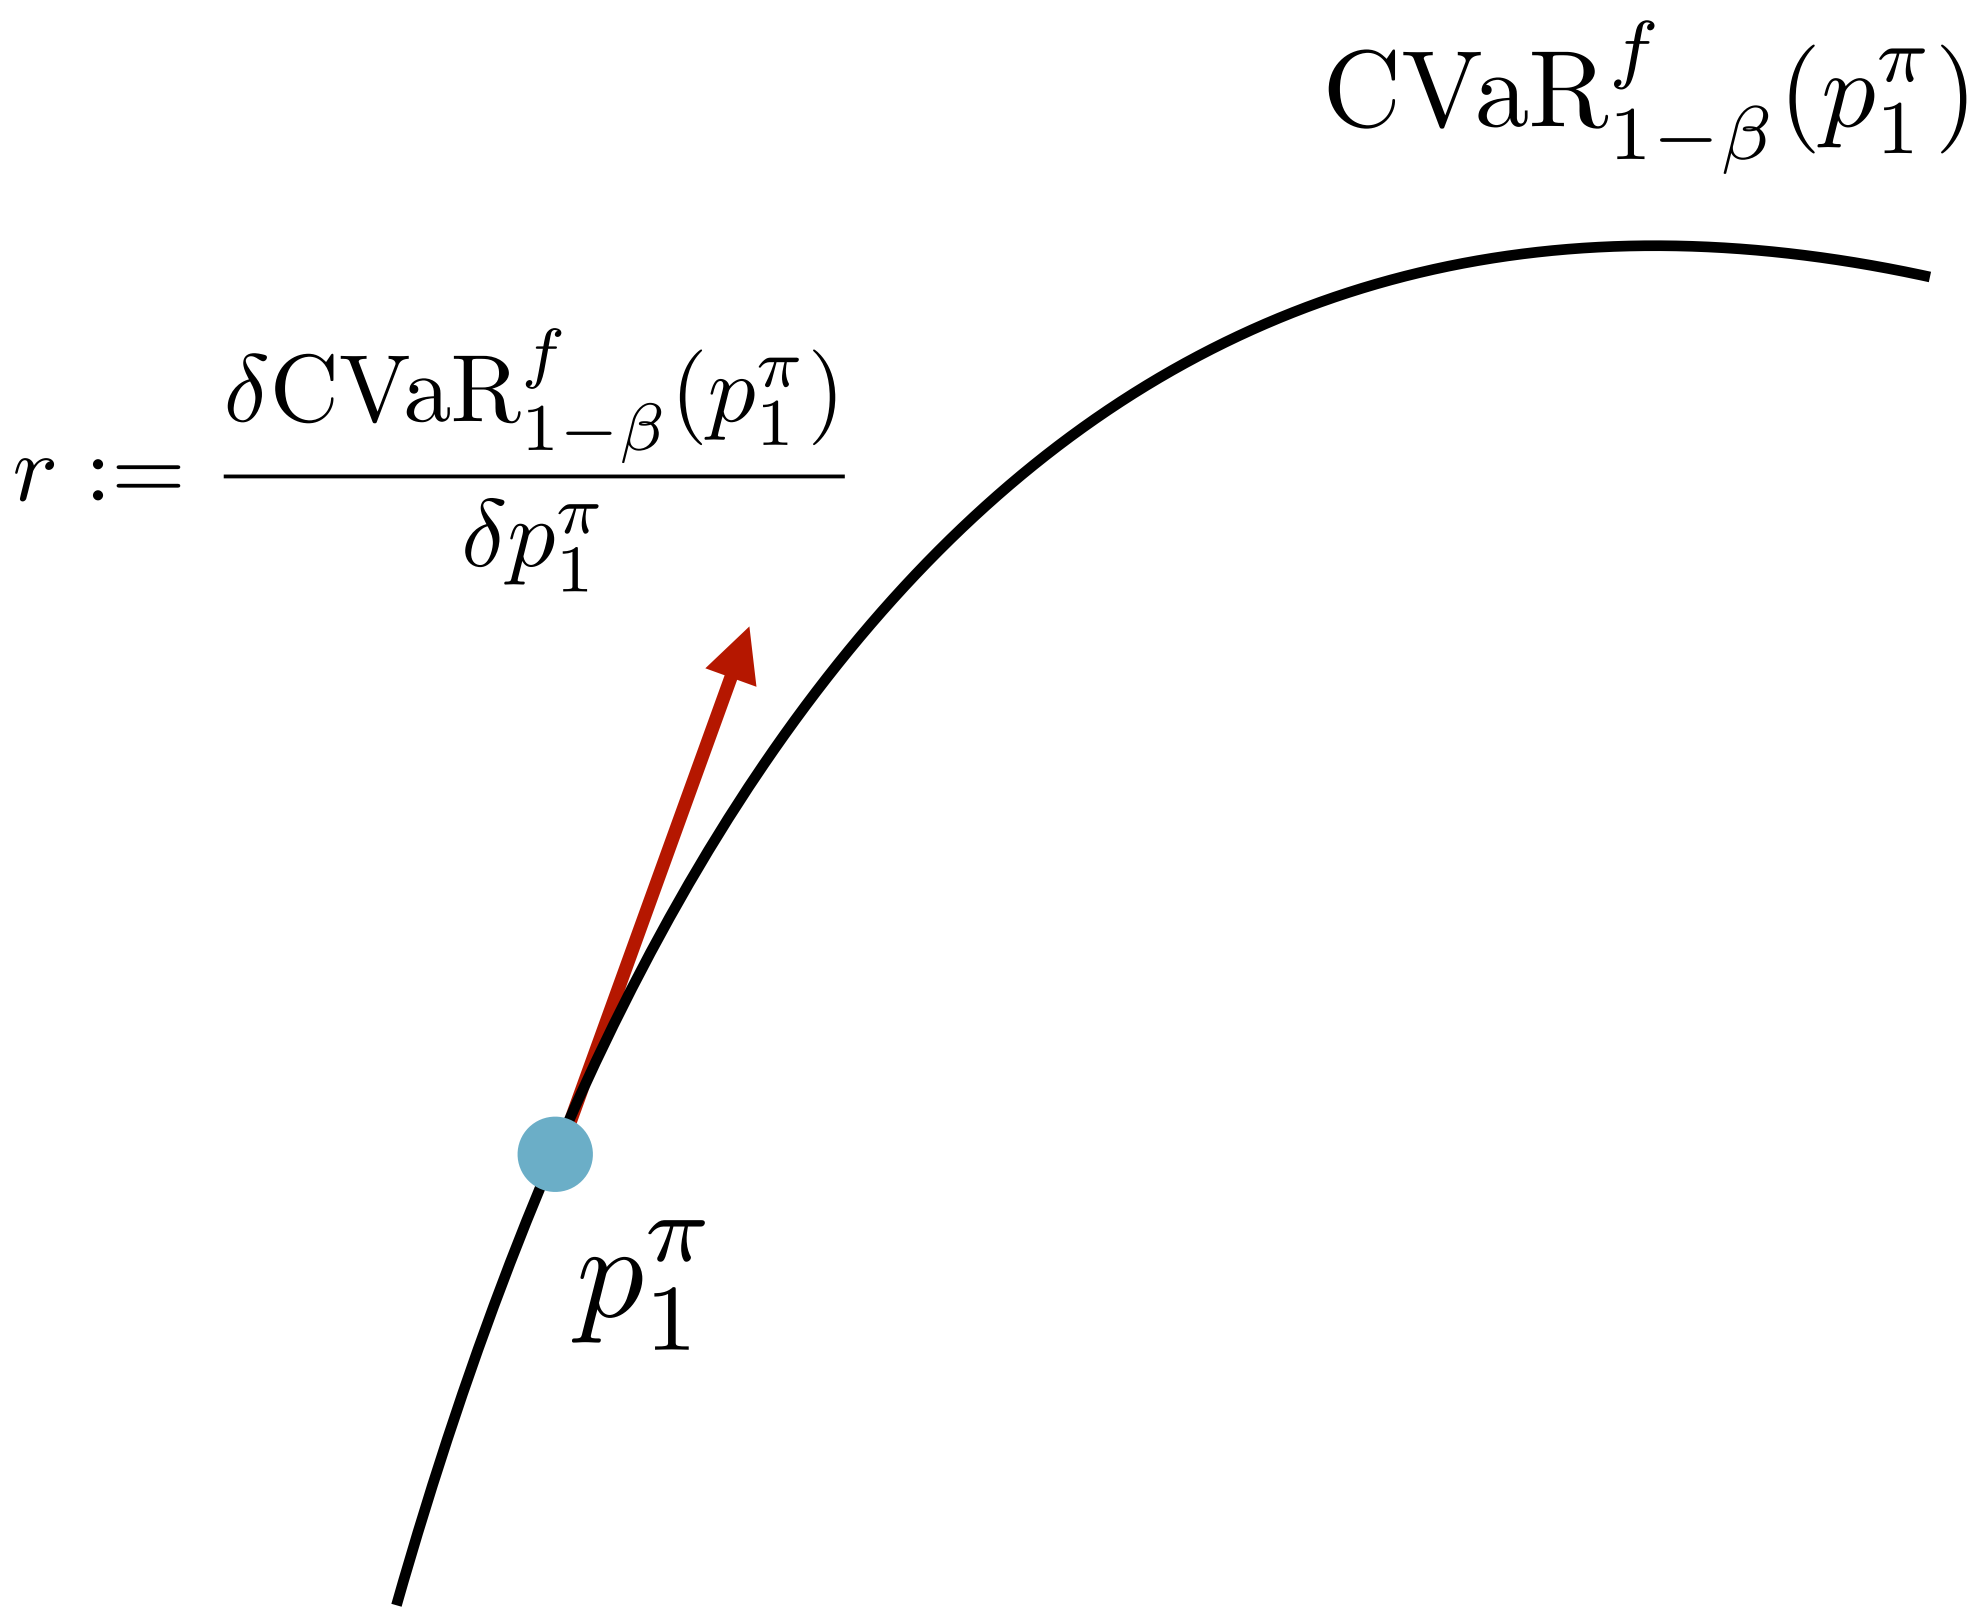

CVaR First Variation

$$r := \frac{\delta \mathrm{CVaR}^f_{1-\beta}(p_1^\pi)}{\delta p_1^\pi} = \frac{(f(x) - VaR^f_{1-\beta}(p_1^\pi))}{\beta}$$

↓

Can use RL fine-tuning

Gradient of CVaR First Variation

$$\nabla r = \frac{\mathbf{1}\{f(x) > VaR^f_{1-\beta}(p_1^\pi)\}}{\beta}$$

↓

Can use gradient-based fine-tuning

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \dots, K$:

$$\mathcal{G} = \mathrm{CVaR}^{f}_{1-\beta} - \alpha \mathbf{KL}$$

$$r_k := \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}}$$

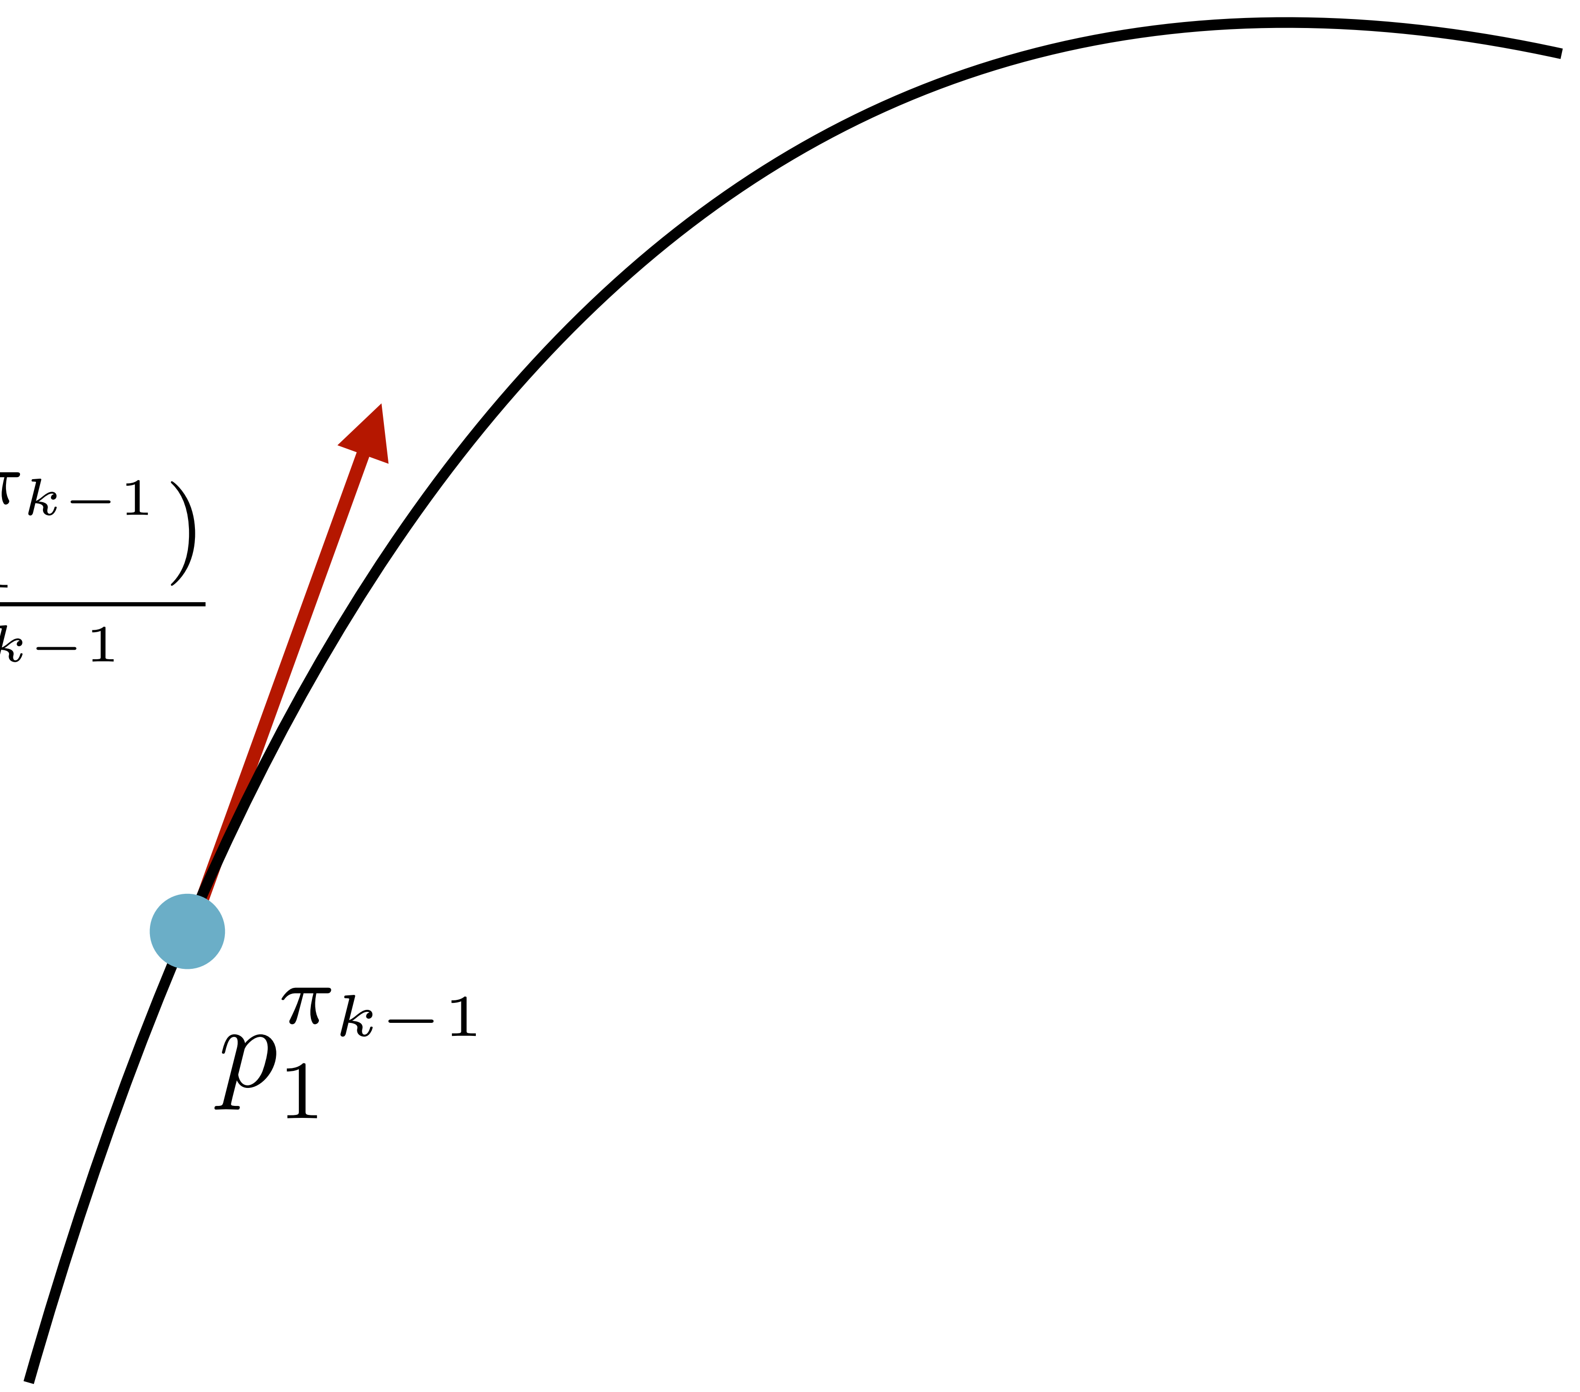

$p_1^{\pi_{k-1}}$

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \ldots, K$:

$$\mathcal{G} = \mathrm{CVaR}^{f}_{1-\beta} - \alpha \mathbf{KL}$$

$$r_k := \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}}$$

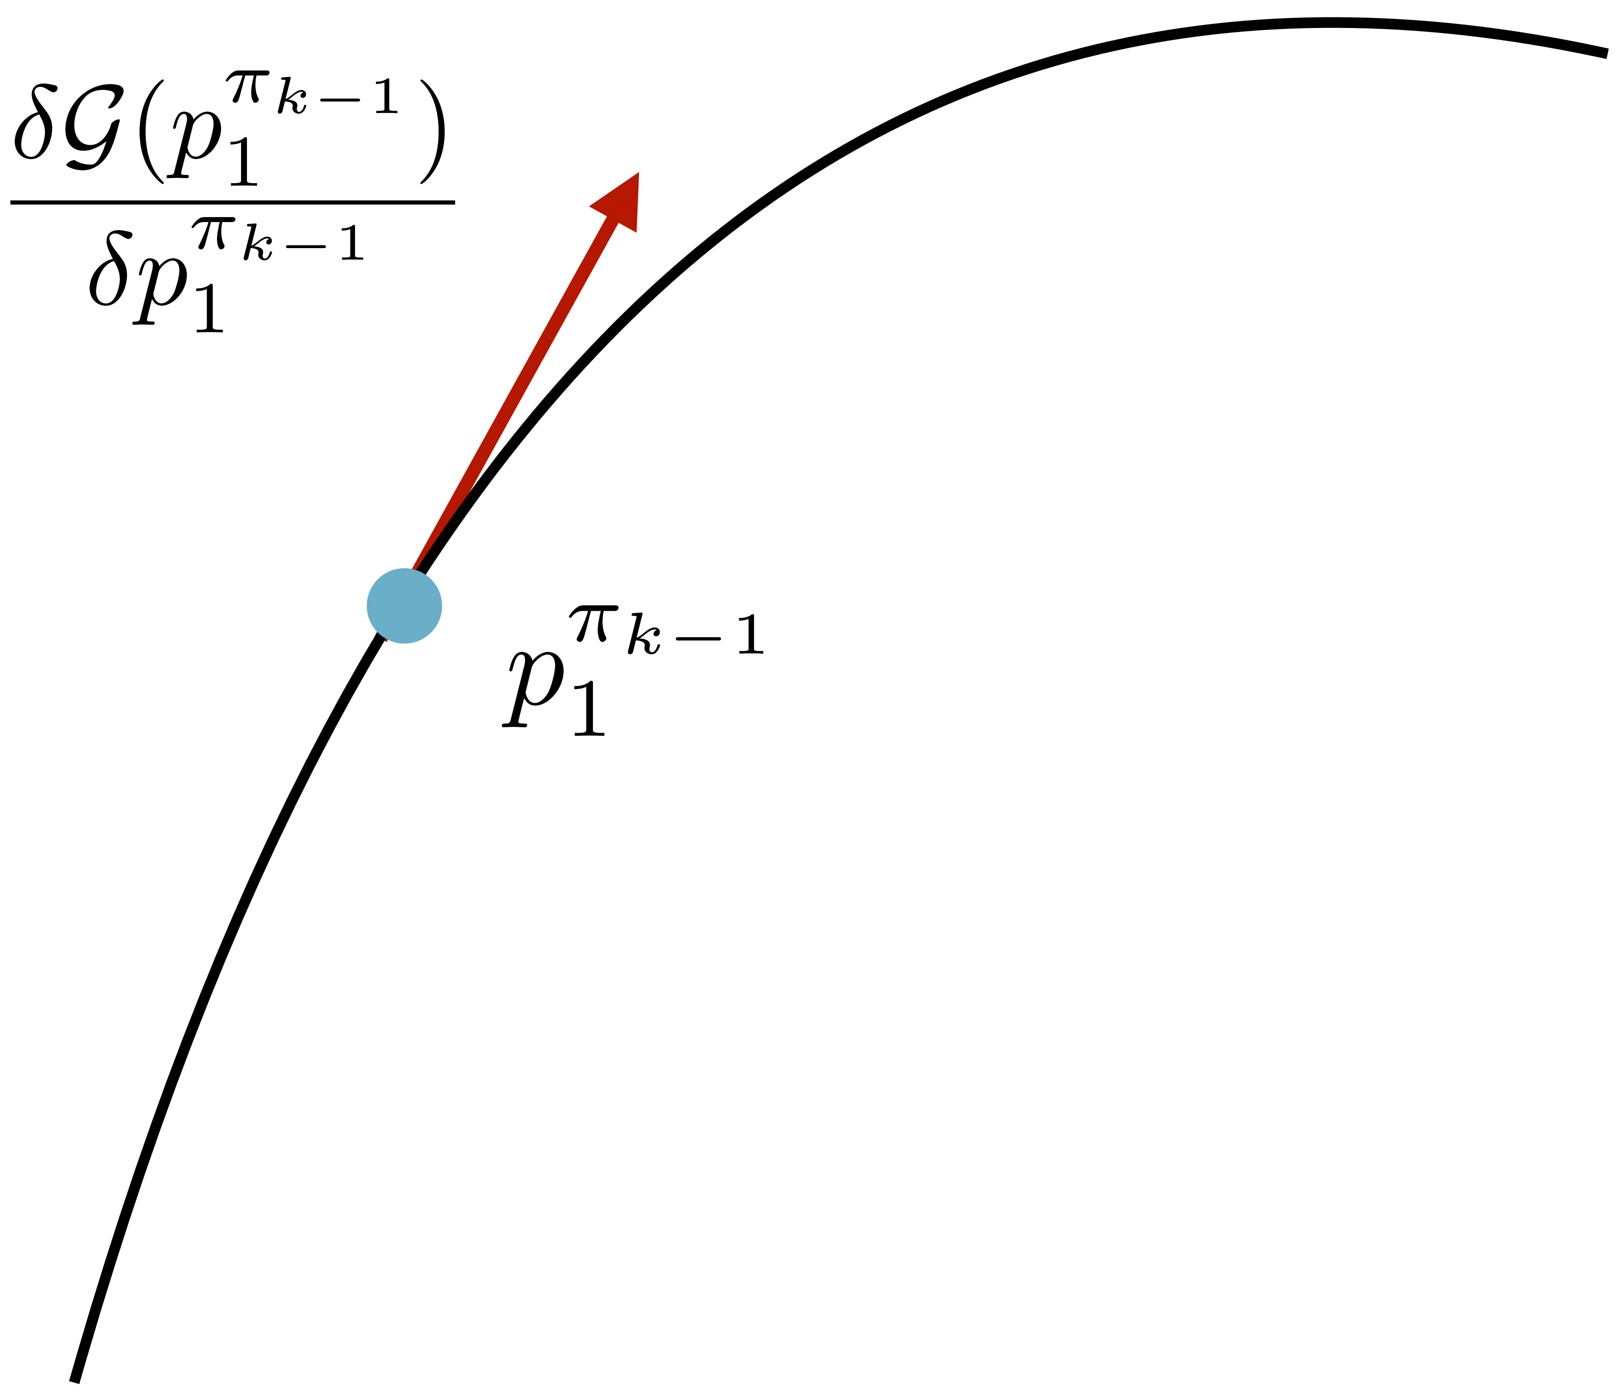

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \ldots, K$:

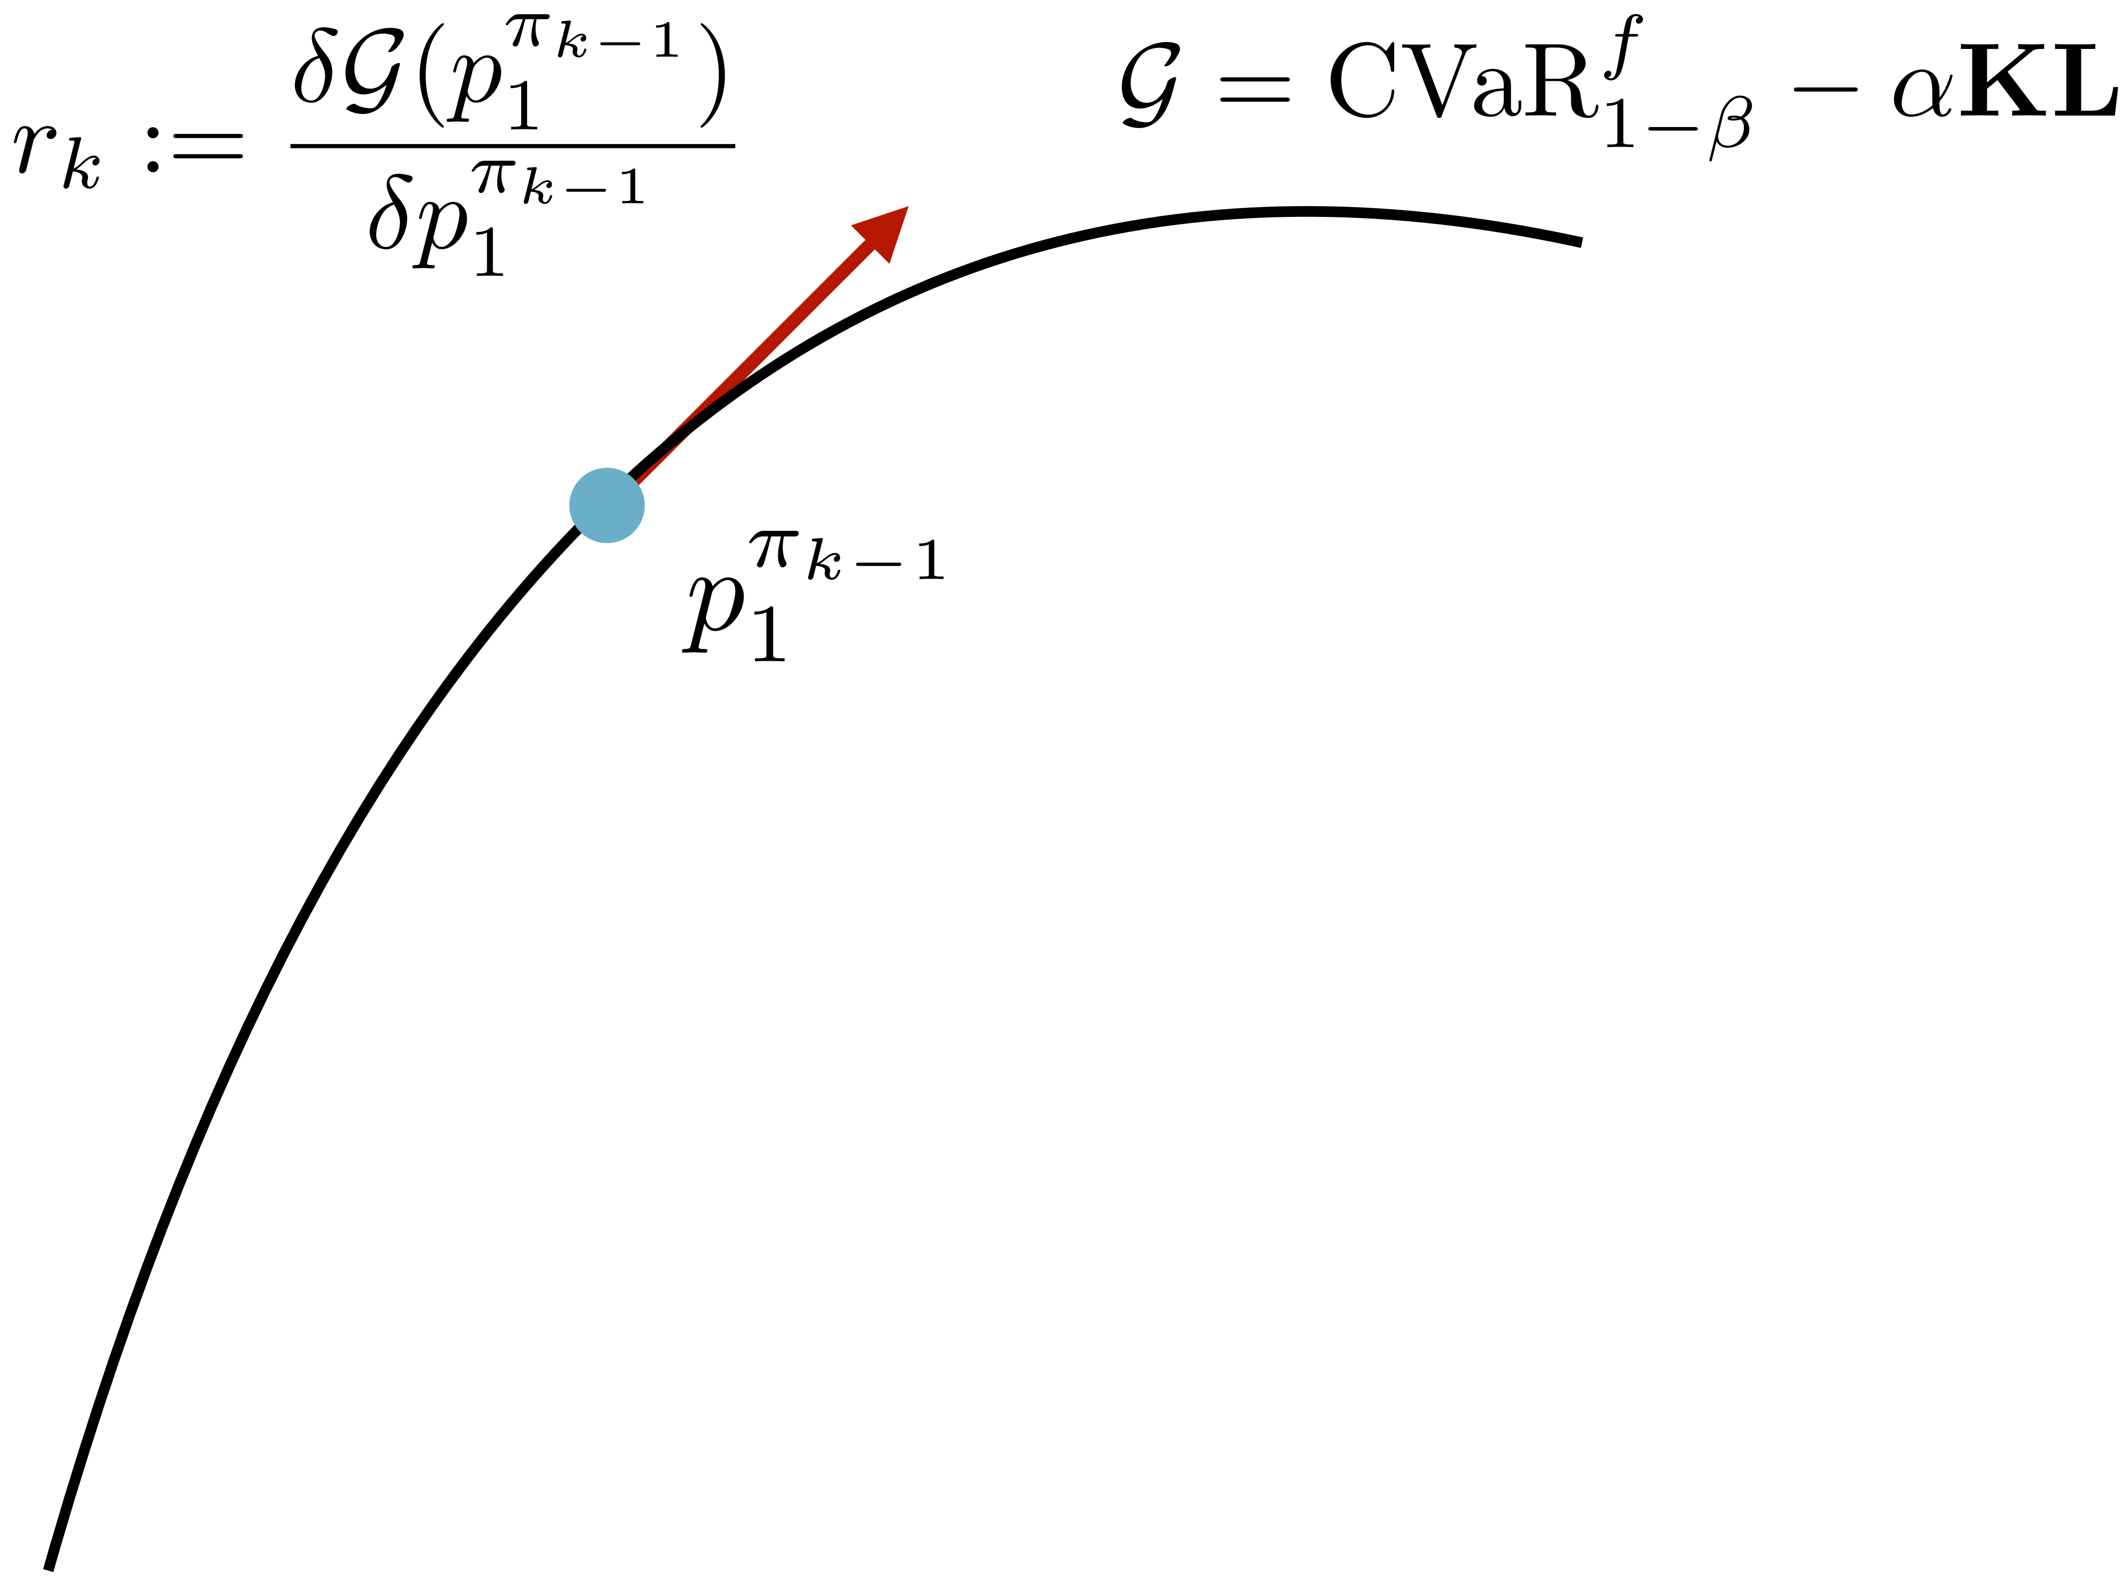

Return $\pi := \pi_K$

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \dots, K$:

Set $\nabla r_k := \nabla \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}} = \frac{\mathbf{1}\{f(x) > VaR_{1-\beta}^{f}(p_k^{\pi})\}}{\beta}$

Fine-tune π_k via standard reward-guided fine-tuning:

$$\pi_k \leftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^{\pi}} [r_k(x)] - \eta_k \mathbf{KL}(p_1^{\pi} \,\|\, p_1^{k-1})$$

Return $\pi := \pi_K$

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \dots, K$:

Set $\nabla r_k := \nabla \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}} = \frac{\mathbf{1}\{f(x) > VaR_{1-\beta}^f(p_k^\pi)\}}{\beta}$

Fine-tune π_k via standard reward-guided fine-tuning:

$$\pi_k \leftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi}[r_k(x)] - \eta_k \mathbf{KL}(p_1^\pi \,\|\, p_1^{k-1})$$

Return $\pi := \pi_K$

Takeaway

1. **CVaR** functional is a **better objective for discovery**
2. **CVaR optimization** can be **reduced to standard RL** → But FDC requires K iterations!

TFFT: dual CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

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

TFFT: dual CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

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

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

CVaR variational formulation

TFFT: dual CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

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

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

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)

TFFT: dual CVaR Fine-Tuning in **One Step**

Tail-Aware Reward-Guided Flow Fine-Tuning

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

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

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:

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

~ 2 lines of code!

CVaR Flow Adaptation for De Novo Molecular Design

GEOM-Drugs

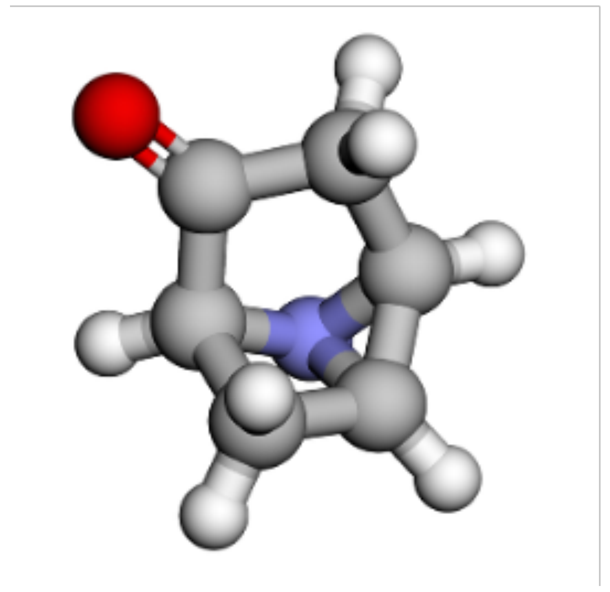

$$g = (\overbrace{X, A, C}, \underbrace{E})$$

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

Edge Features
- Bond Type (D)

Flow over 3D molecular conformers

[Dunn and Koes, 2024]

CVaR Flow Adaptation for De Novo Molecular Design

GEOM-Drugs

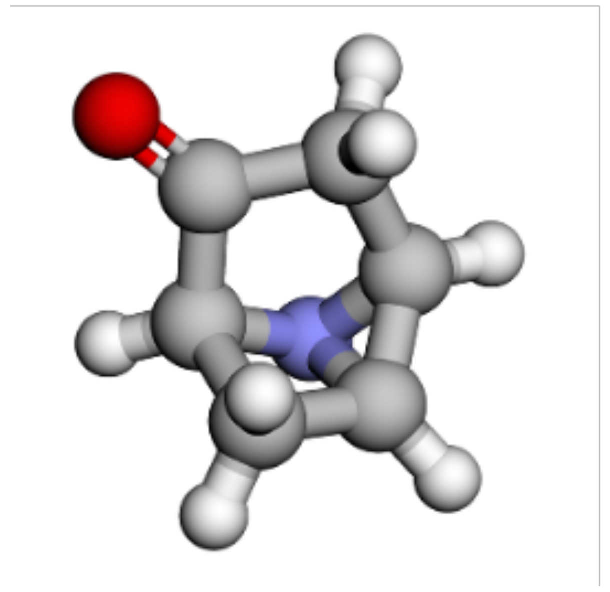

$$g = (X, A, C, E)$$

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

Edge Features
- Bond Type (D)

Flow over 3D molecular conformers

	Training time	FT calls	$\mathbb{E}[r]$	R-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	$\mathbf{29.8}_{\pm 0.58}$	$167.1_{\pm 7.22}$	$78.3\%_{\pm 1.01\%}$	$7.66_{\pm 0.02}$
FDC (K=3)	$\approx$ 3T	3	$29.7_{\pm 0.85}$	$172.5_{\pm 11.8}$	$82.7\%_{\pm 0.47\%}$	$\mathbf{7.56}_{\pm 0.02}$
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]

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

Tail-Aware Fine-Tuning for Molecular Design: GEOM-Drugs

Pre-trained

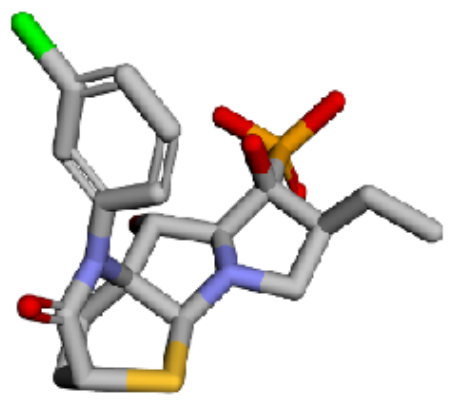

840

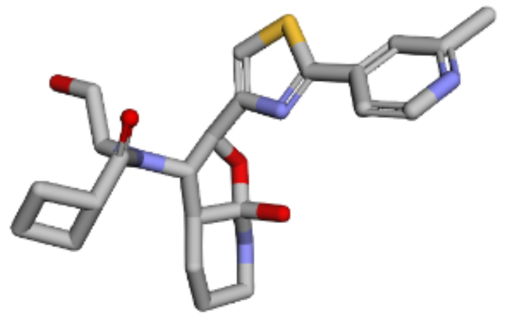

826

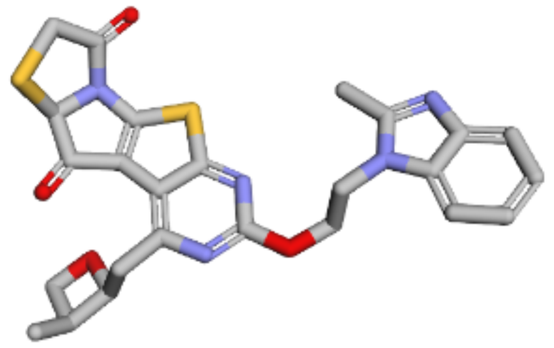

755

Tail-Aware Fine-Tuning for Molecular Design: GEOM-Drugs

Pre-trained

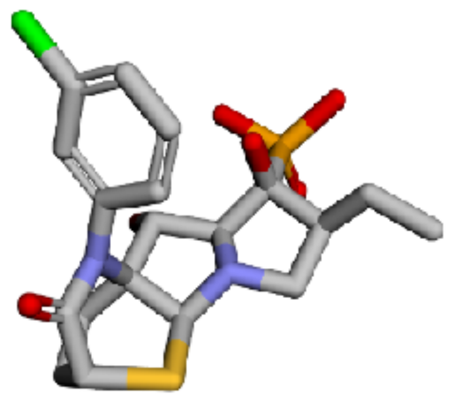

840

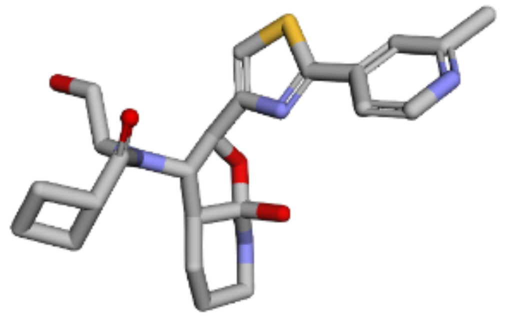

826

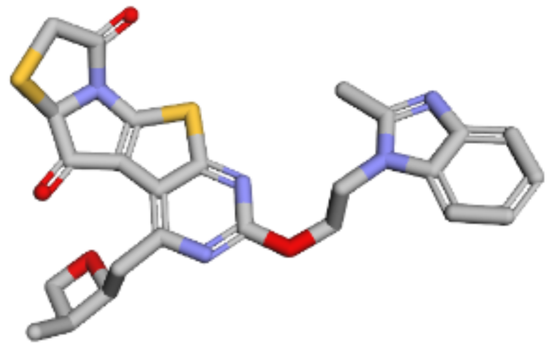

755

Adjoint Matching

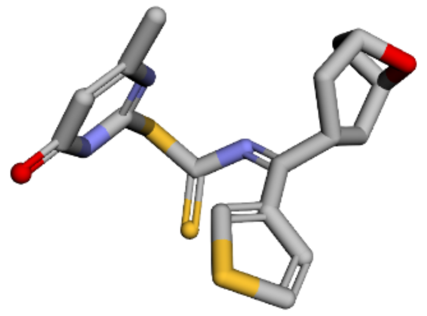

1069

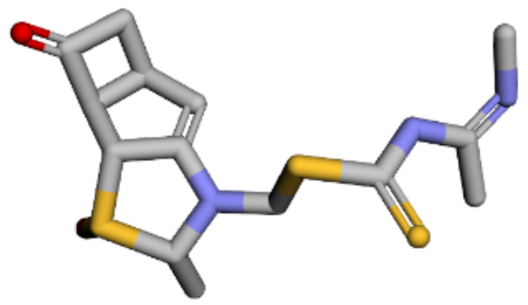

891

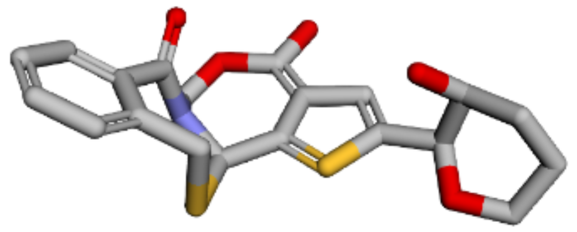

852

Tail-Aware Fine-Tuning for Molecular Design: GEOM-Drugs

Pre-trained

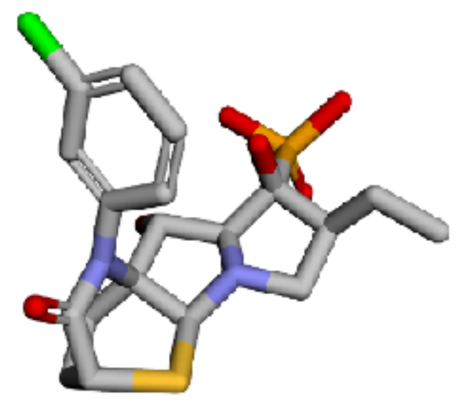

840

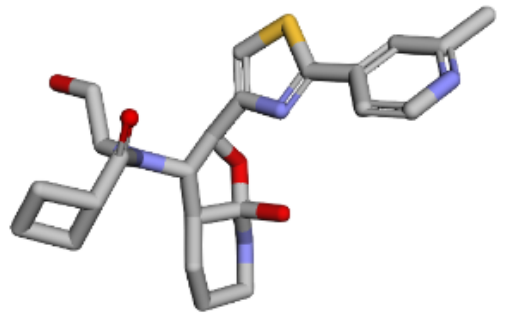

826

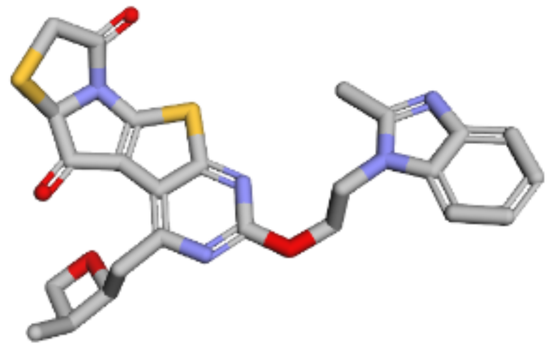

755

Adjoint Matching

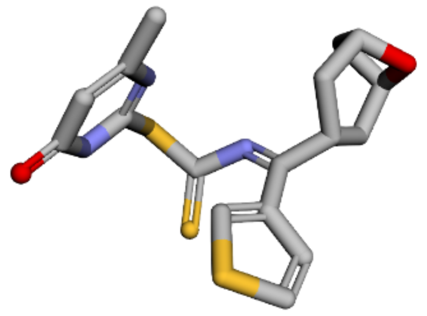

1069

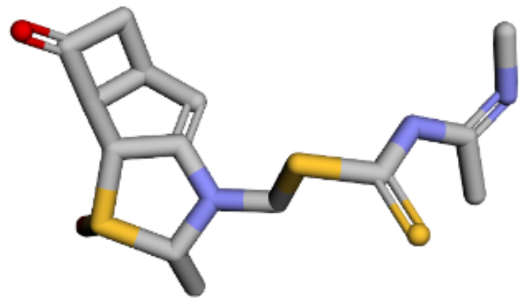

891

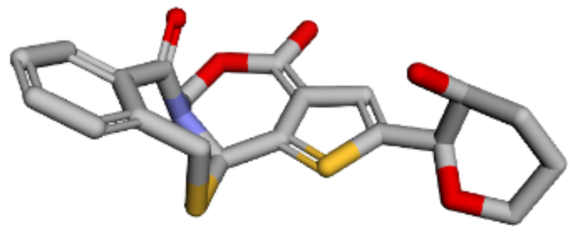

852

FDC

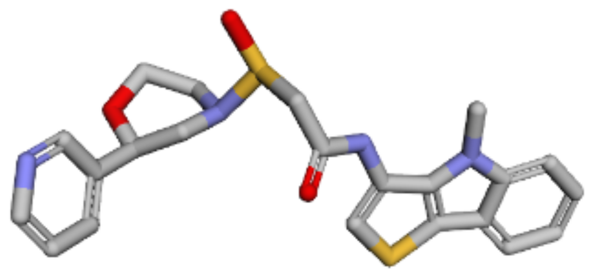

2012

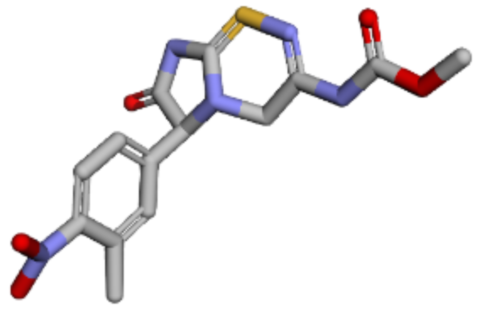

799

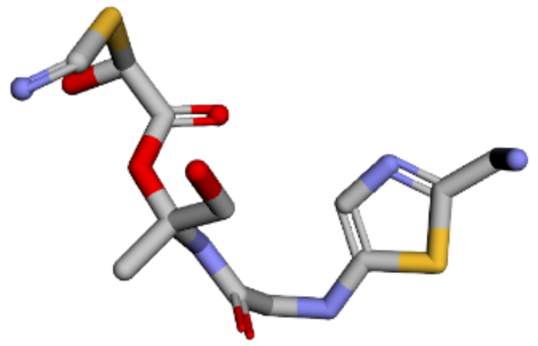

757

~2x
higher-rewards

Tail-Aware Fine-Tuning for Molecular Design: GEOM-Drugs

Pre-trained

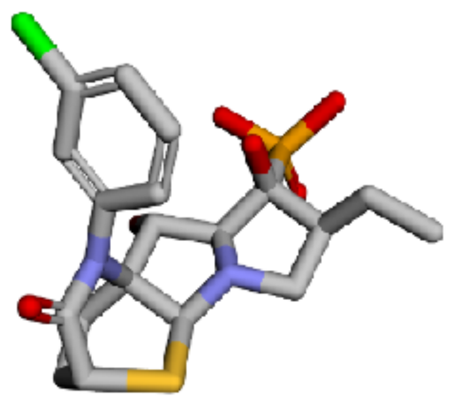

840

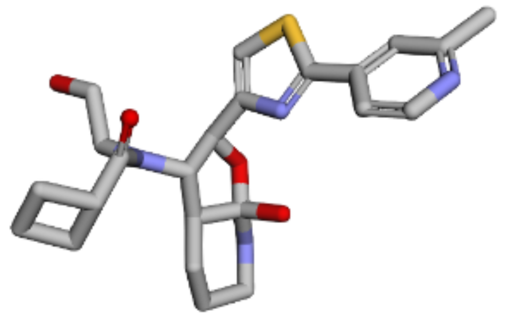

826

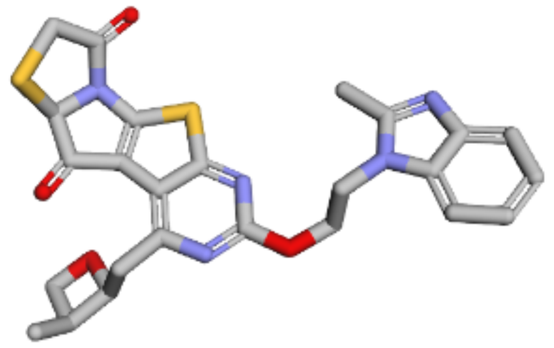

755

Adjoint Matching

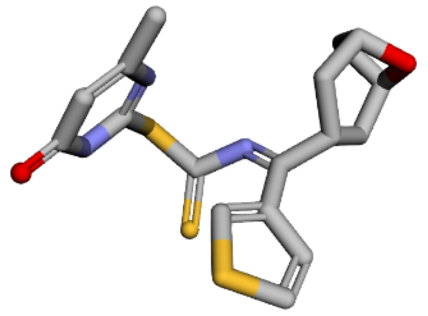

1069

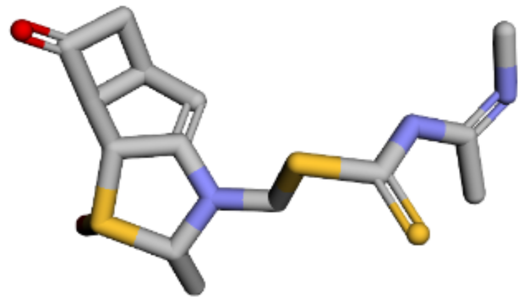

891

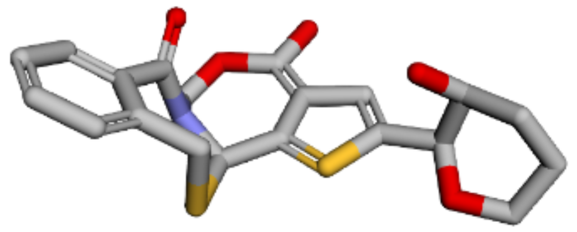

852

FDC

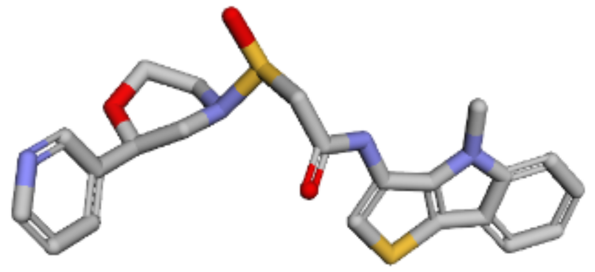

2012

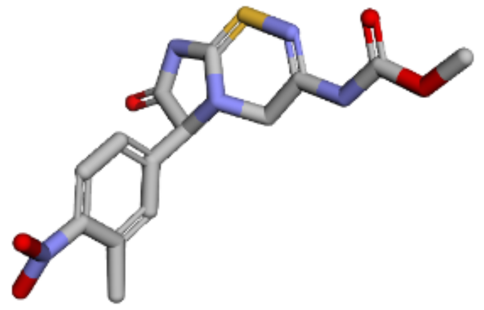

799

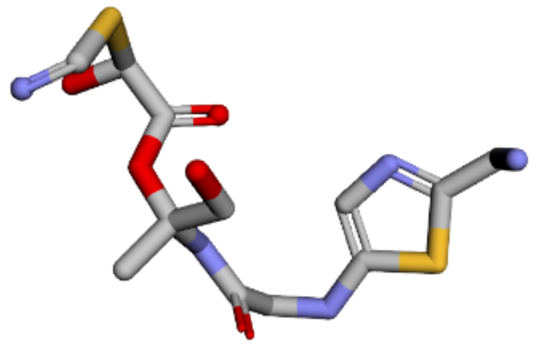

757

TFFT

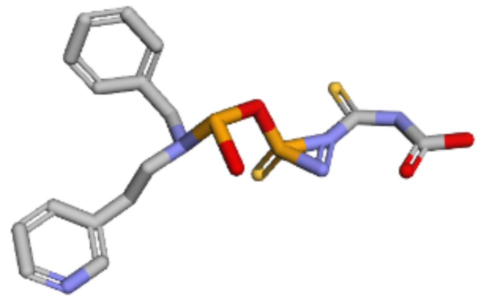

5059

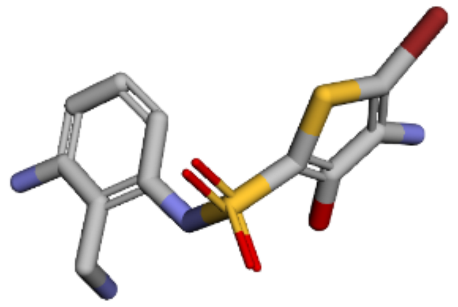

3237

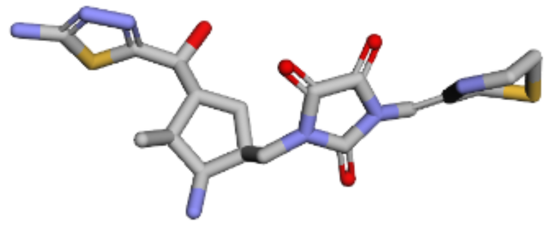

1032

~5x
higher-rewards

This talk:

Foundations of Generative Discovery Beyond the Data

Part I

Tail-Aware Reward Adaptation

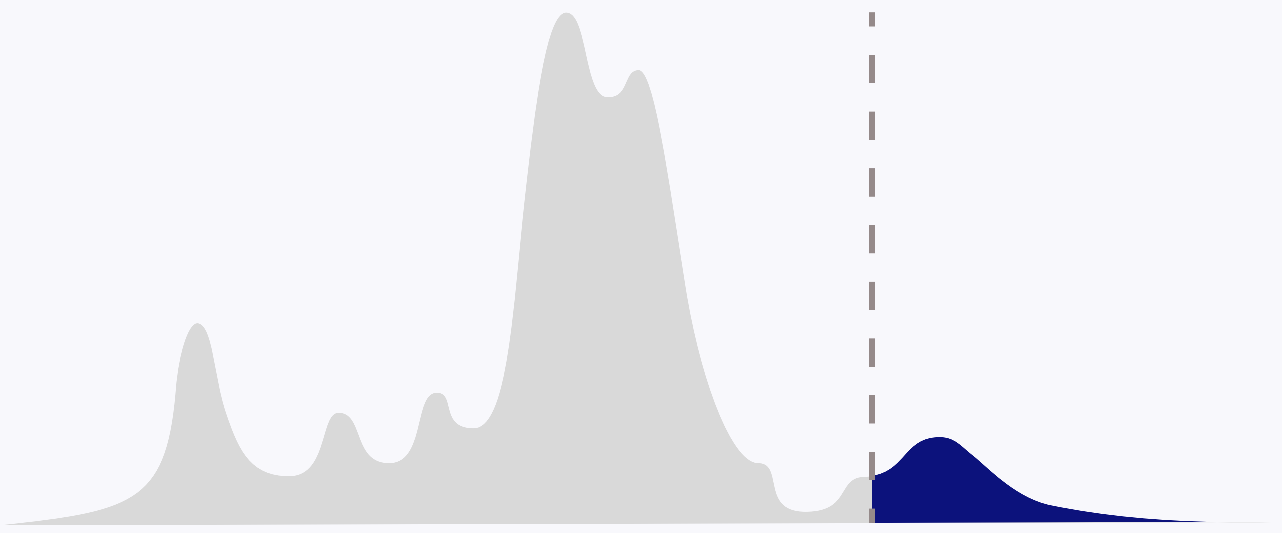

$\text{CVaR}^f_{1-\beta}(p^\pi_1)$

Part II

Data Debiasing and Hidden Mode Discovery

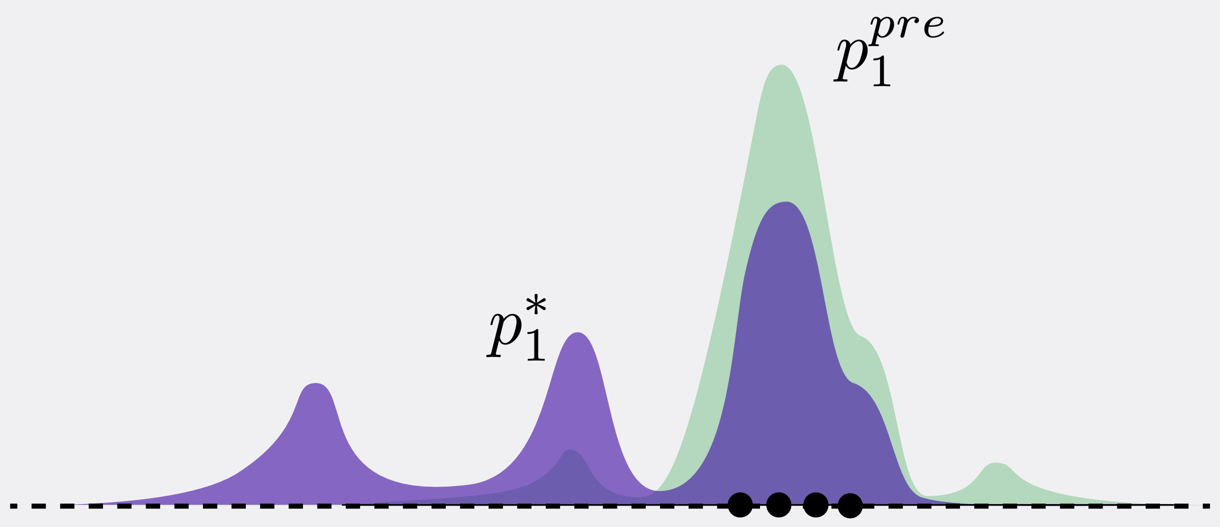

Part III

Out-of-Distribution Flow Modeling

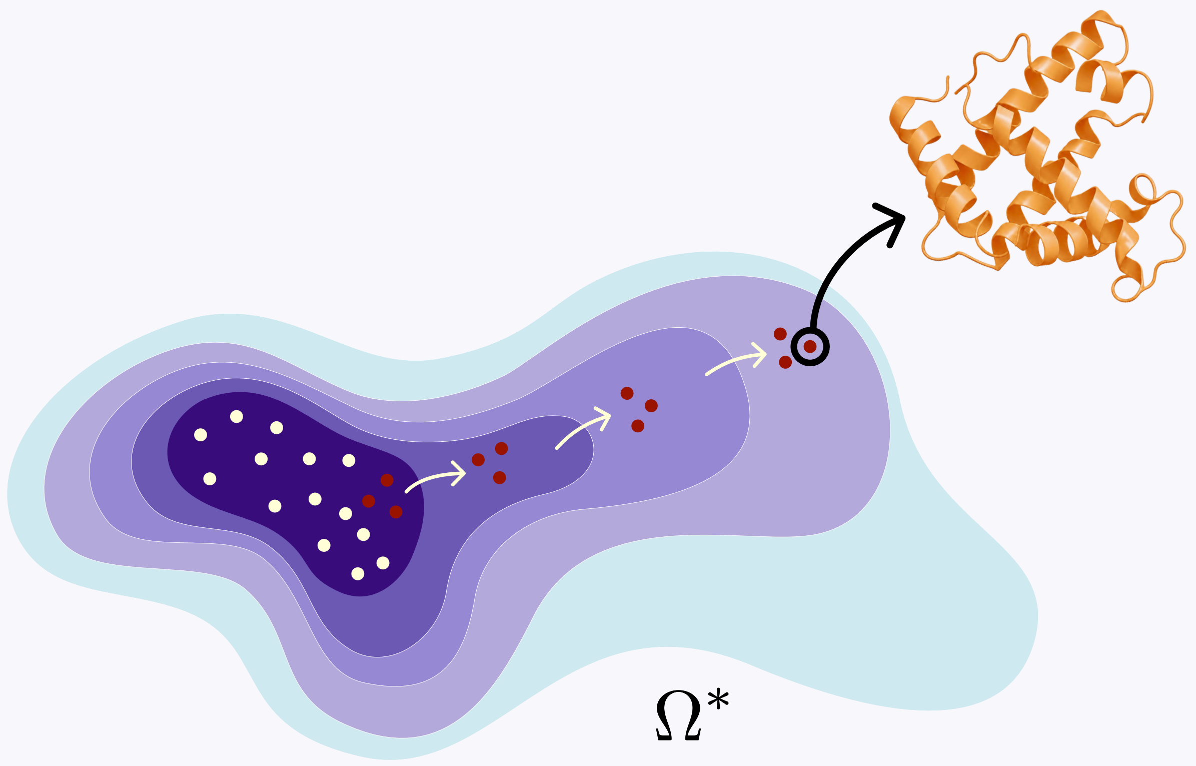

Exploration via Flow Adaptation: Key References

The Importance of Non-Markovianity in Maximum State Entropy Exploration

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

Outstanding Paper Award at **ICML 2022**

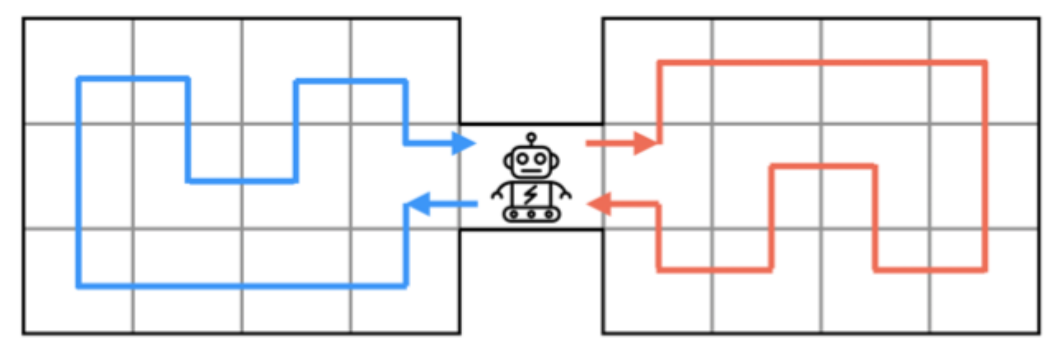

Exploration via Flow Adaptation: Key References

Provable Maximum Entropy Manifold Exploration via Diffusion Models

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

ICML 2025

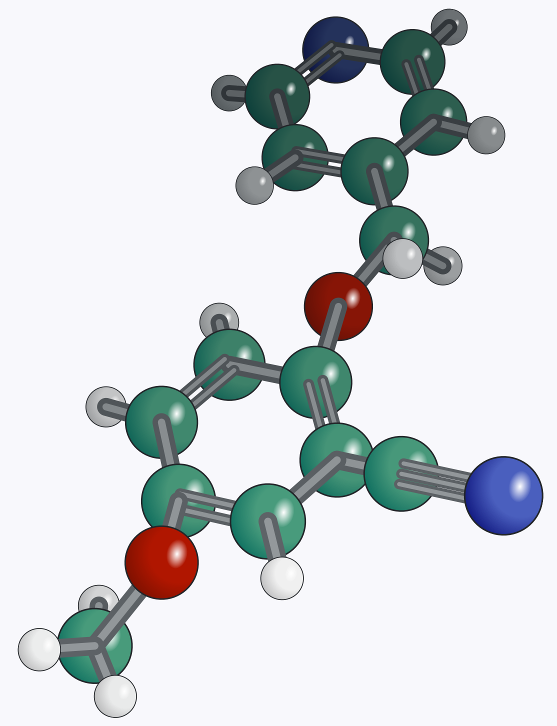

Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning

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

Spotlight NeurIPS 2025   Oral at ICML 2025 BioGen Workshop

The Importance of Non-Markovianity in Maximum State Entropy Exploration

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

Outstanding Paper Award at ICML 2022

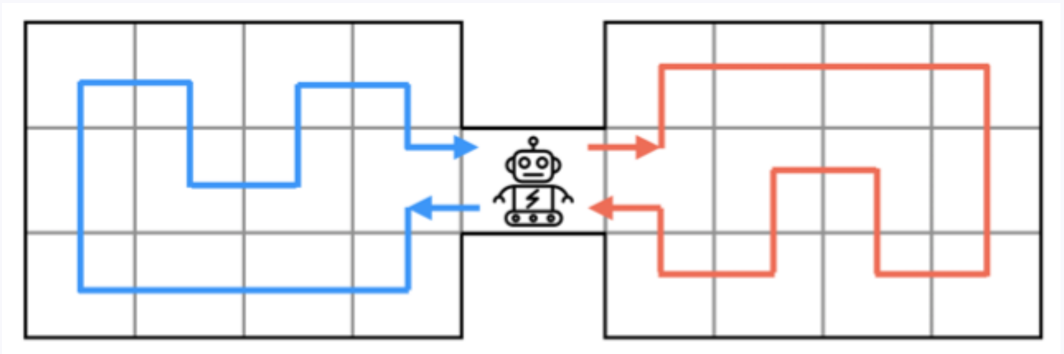

Generative Modeling vs Discovery

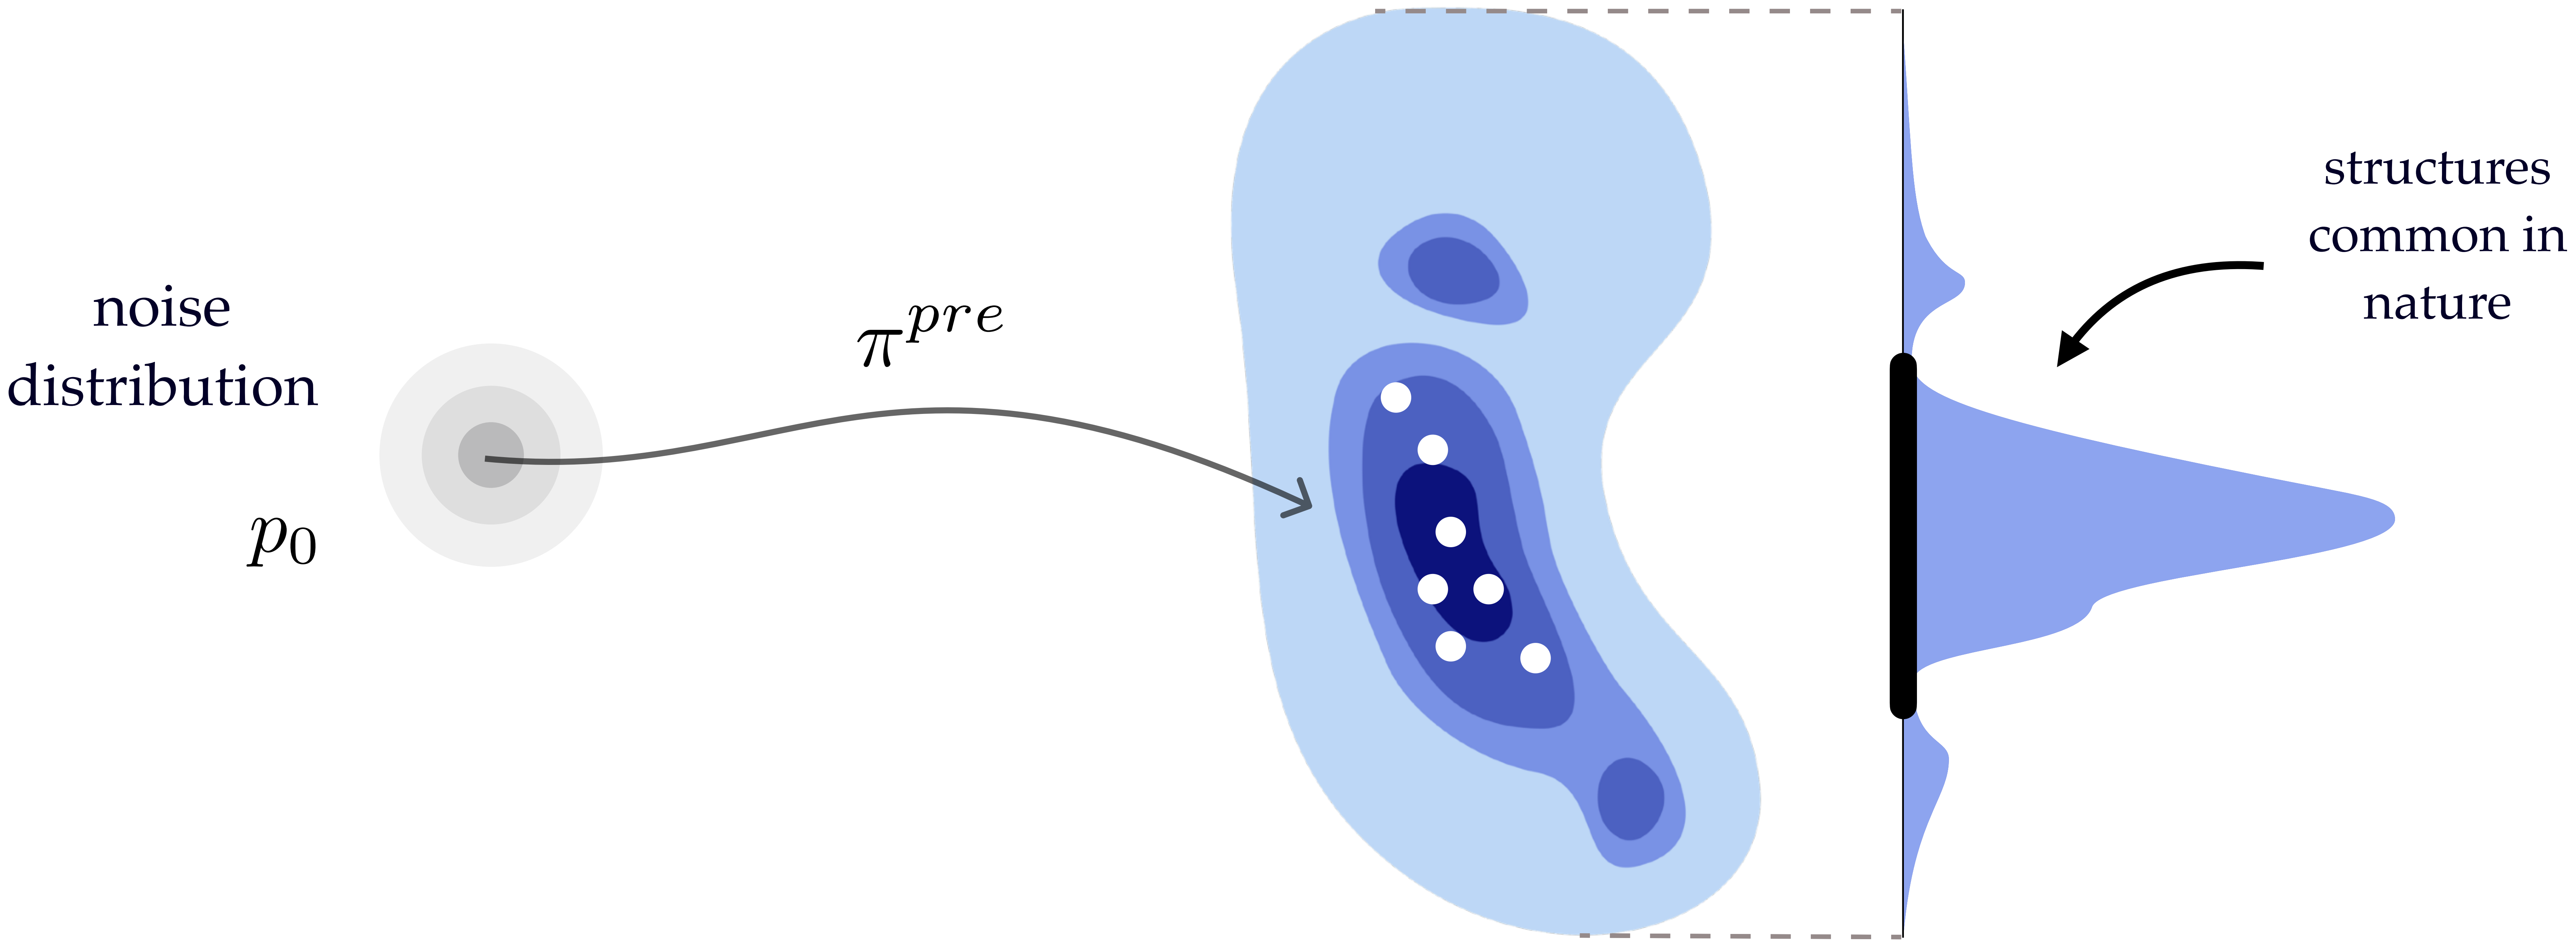

Generative modeling priors approximately **match the frequency of available structures**

Generative Modeling vs Discovery

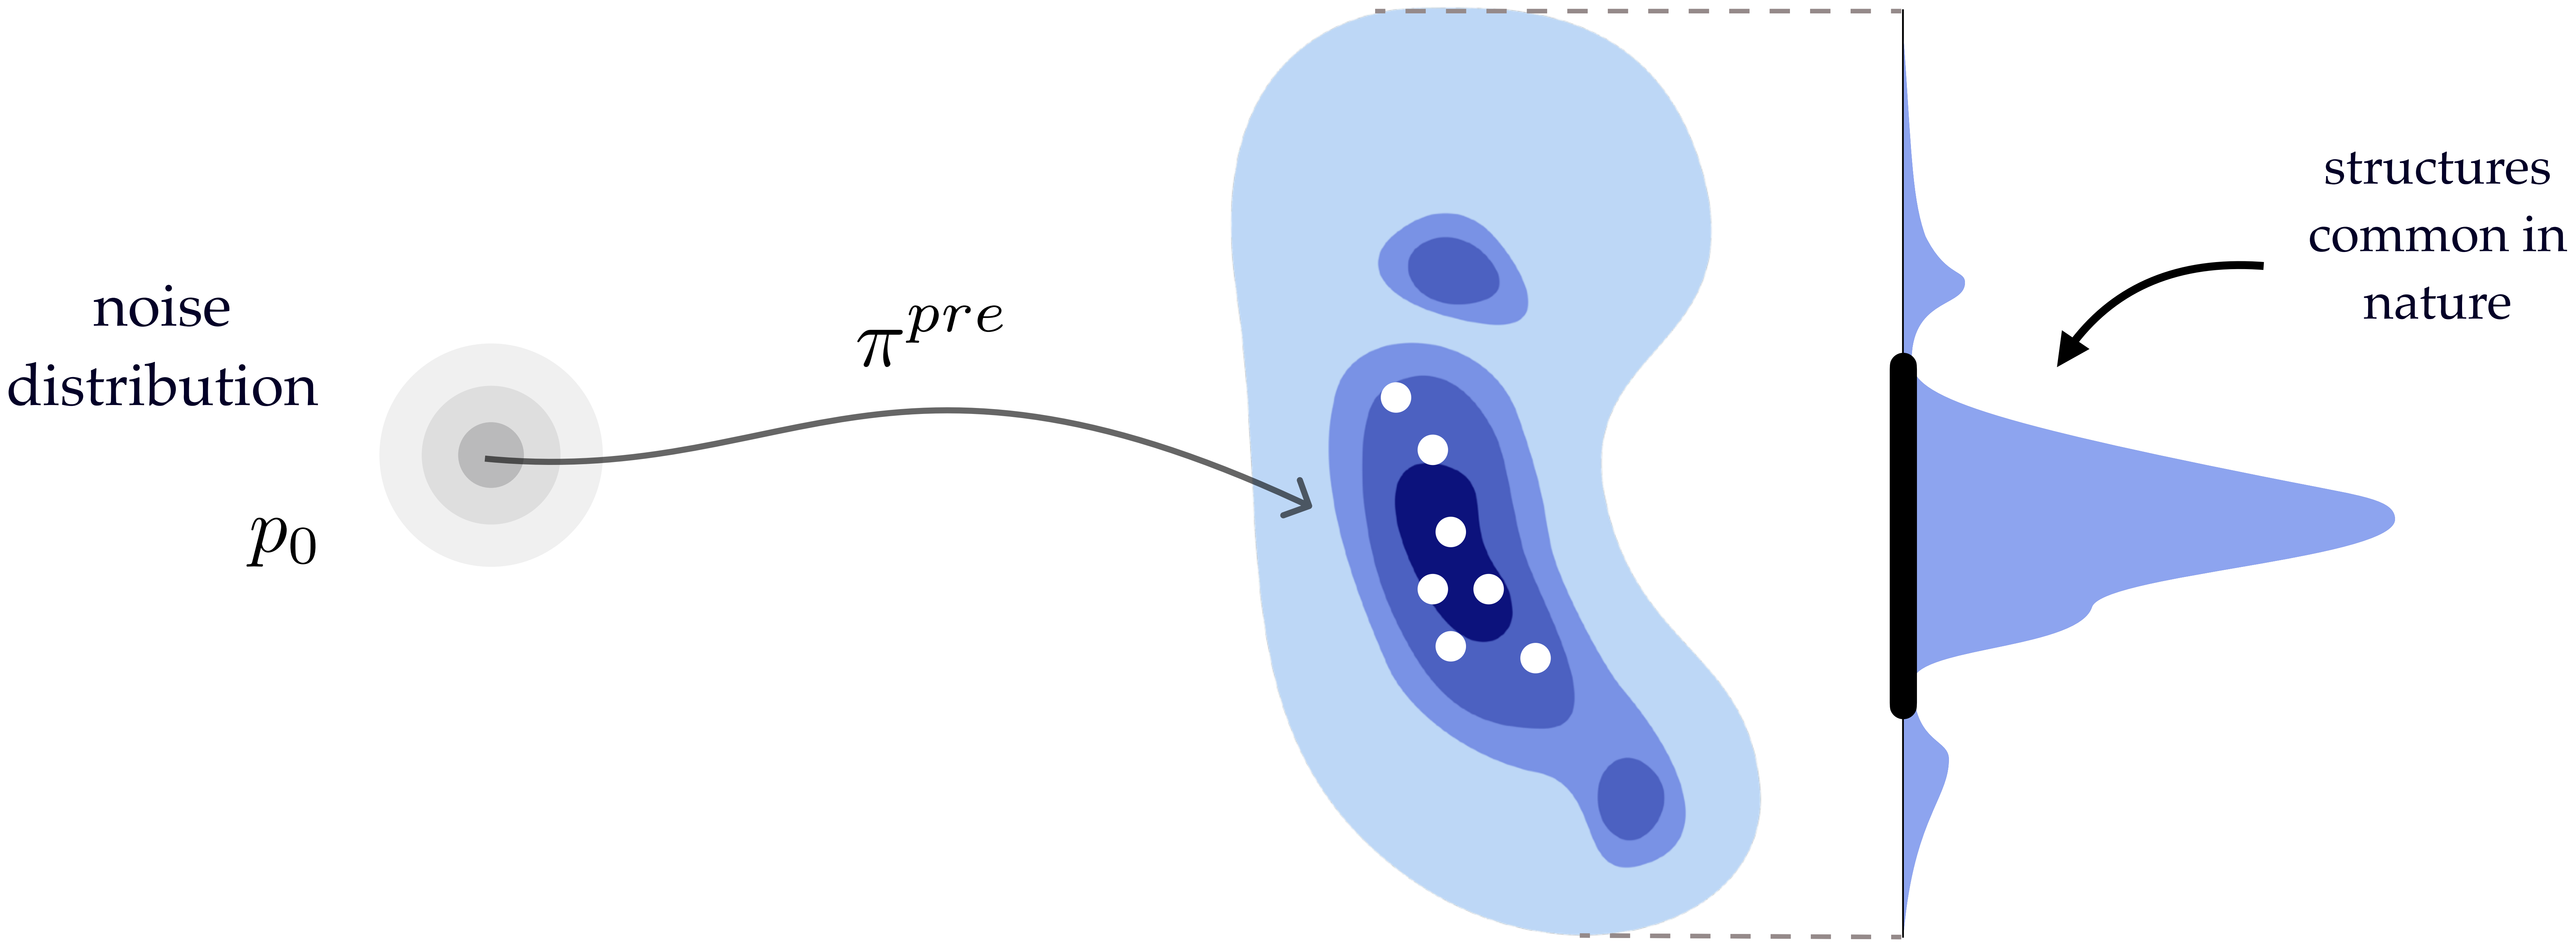

Proteins —X-ray Cryst.→ **PDB** (Protein Data Bank) —Generative modeling→ **AlphaFold**

Models trained on PDB are **collapsed** to static structures.

Generative Modeling vs Discovery

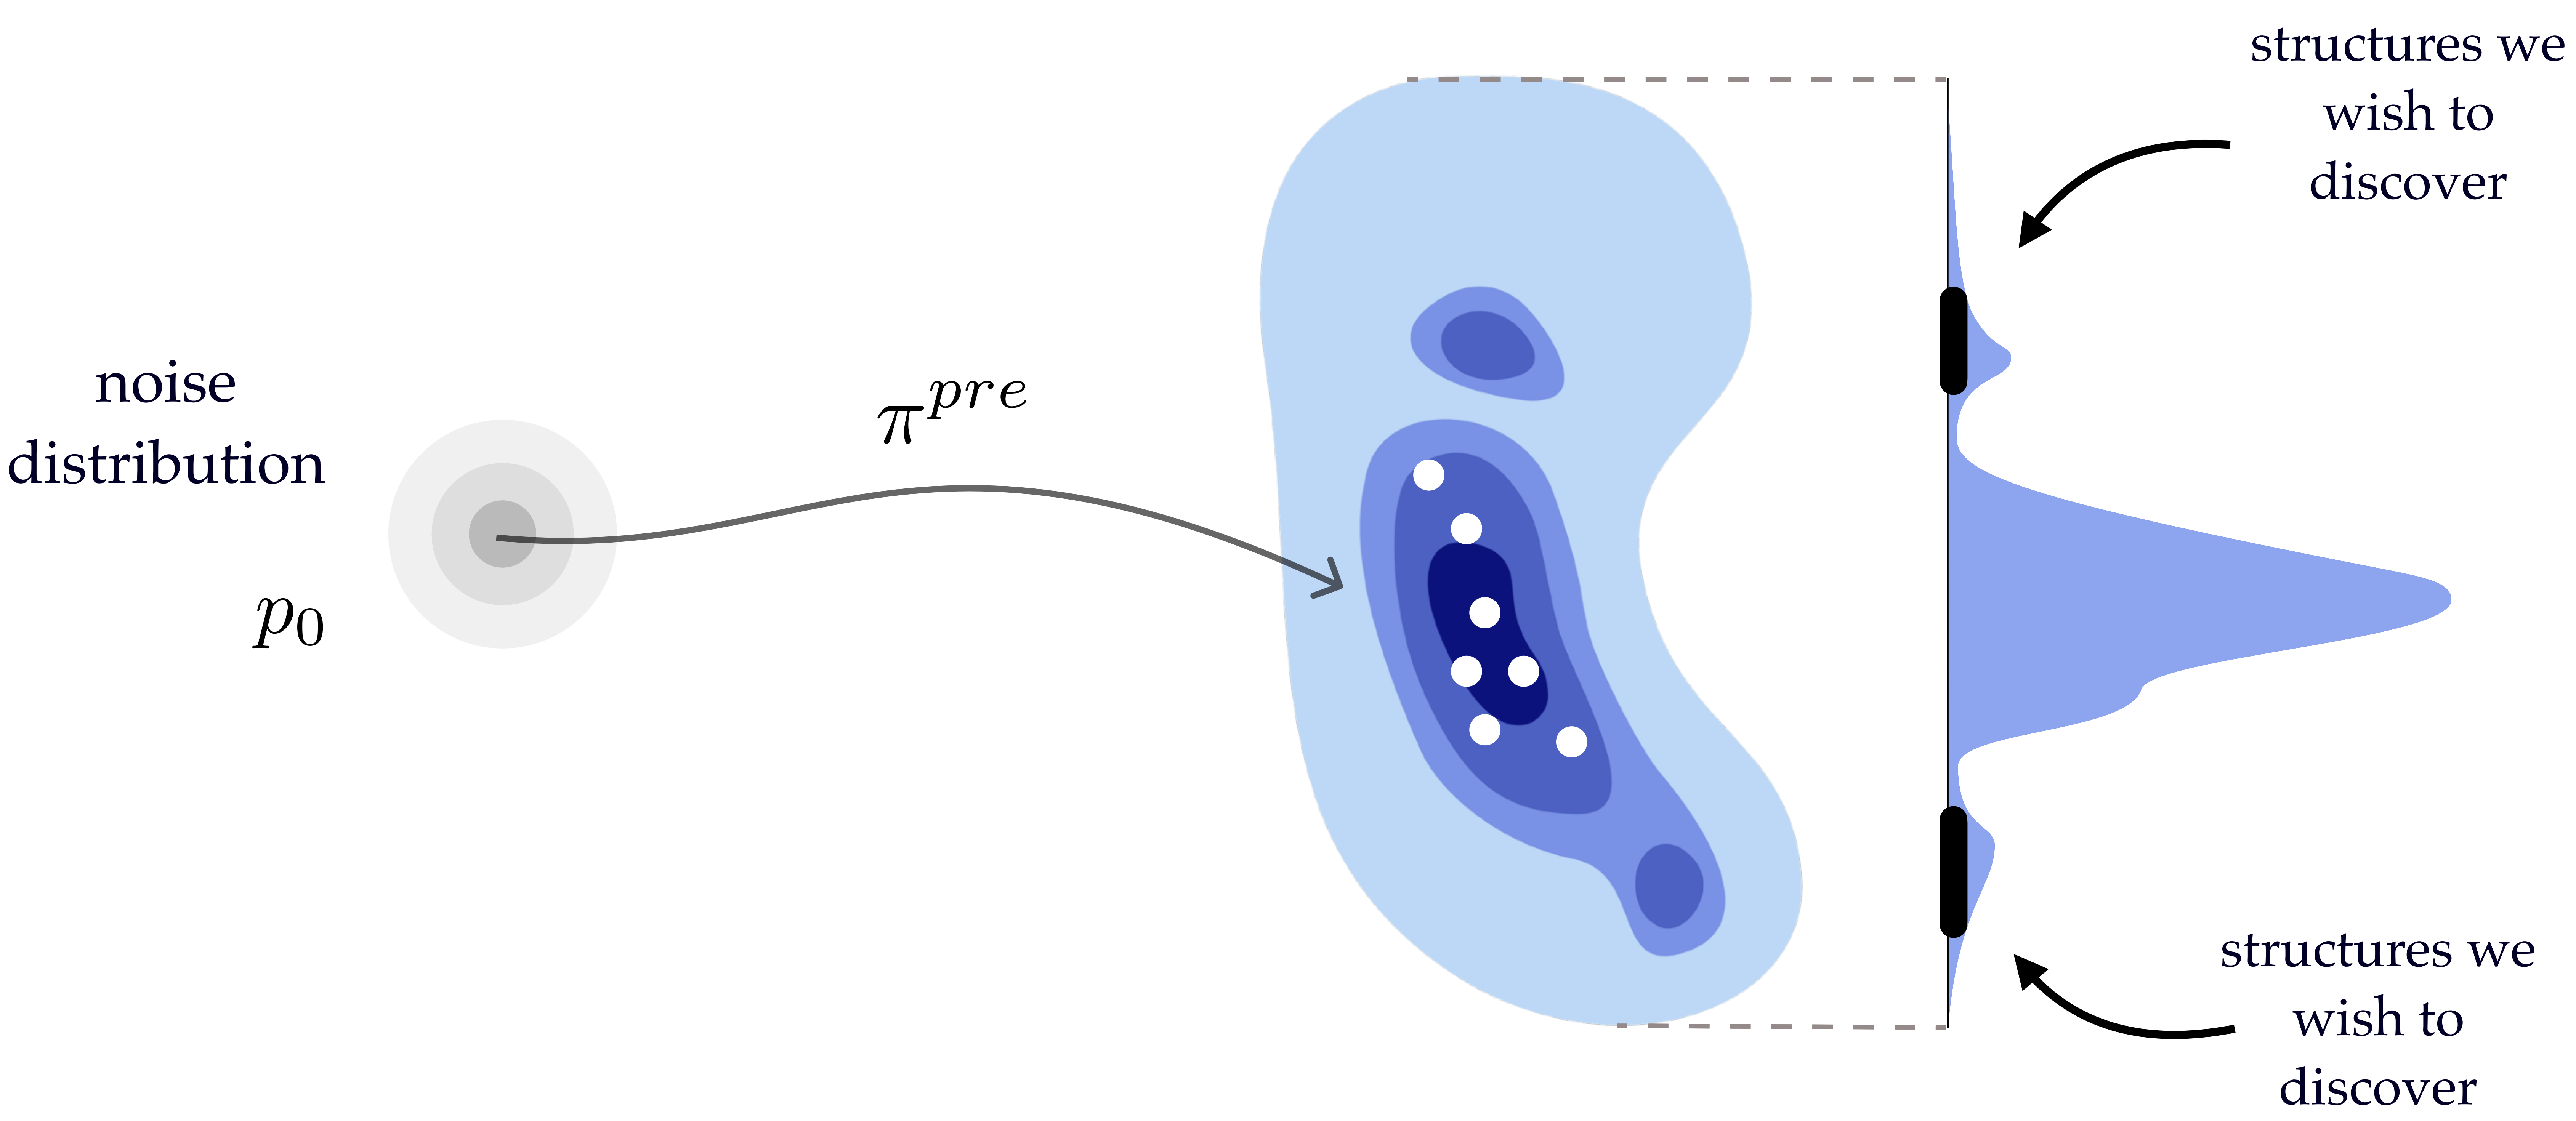

Discovery requires to generate **rare, new structures**

Data Debiasing leads to *Hidden* Mode Discovery

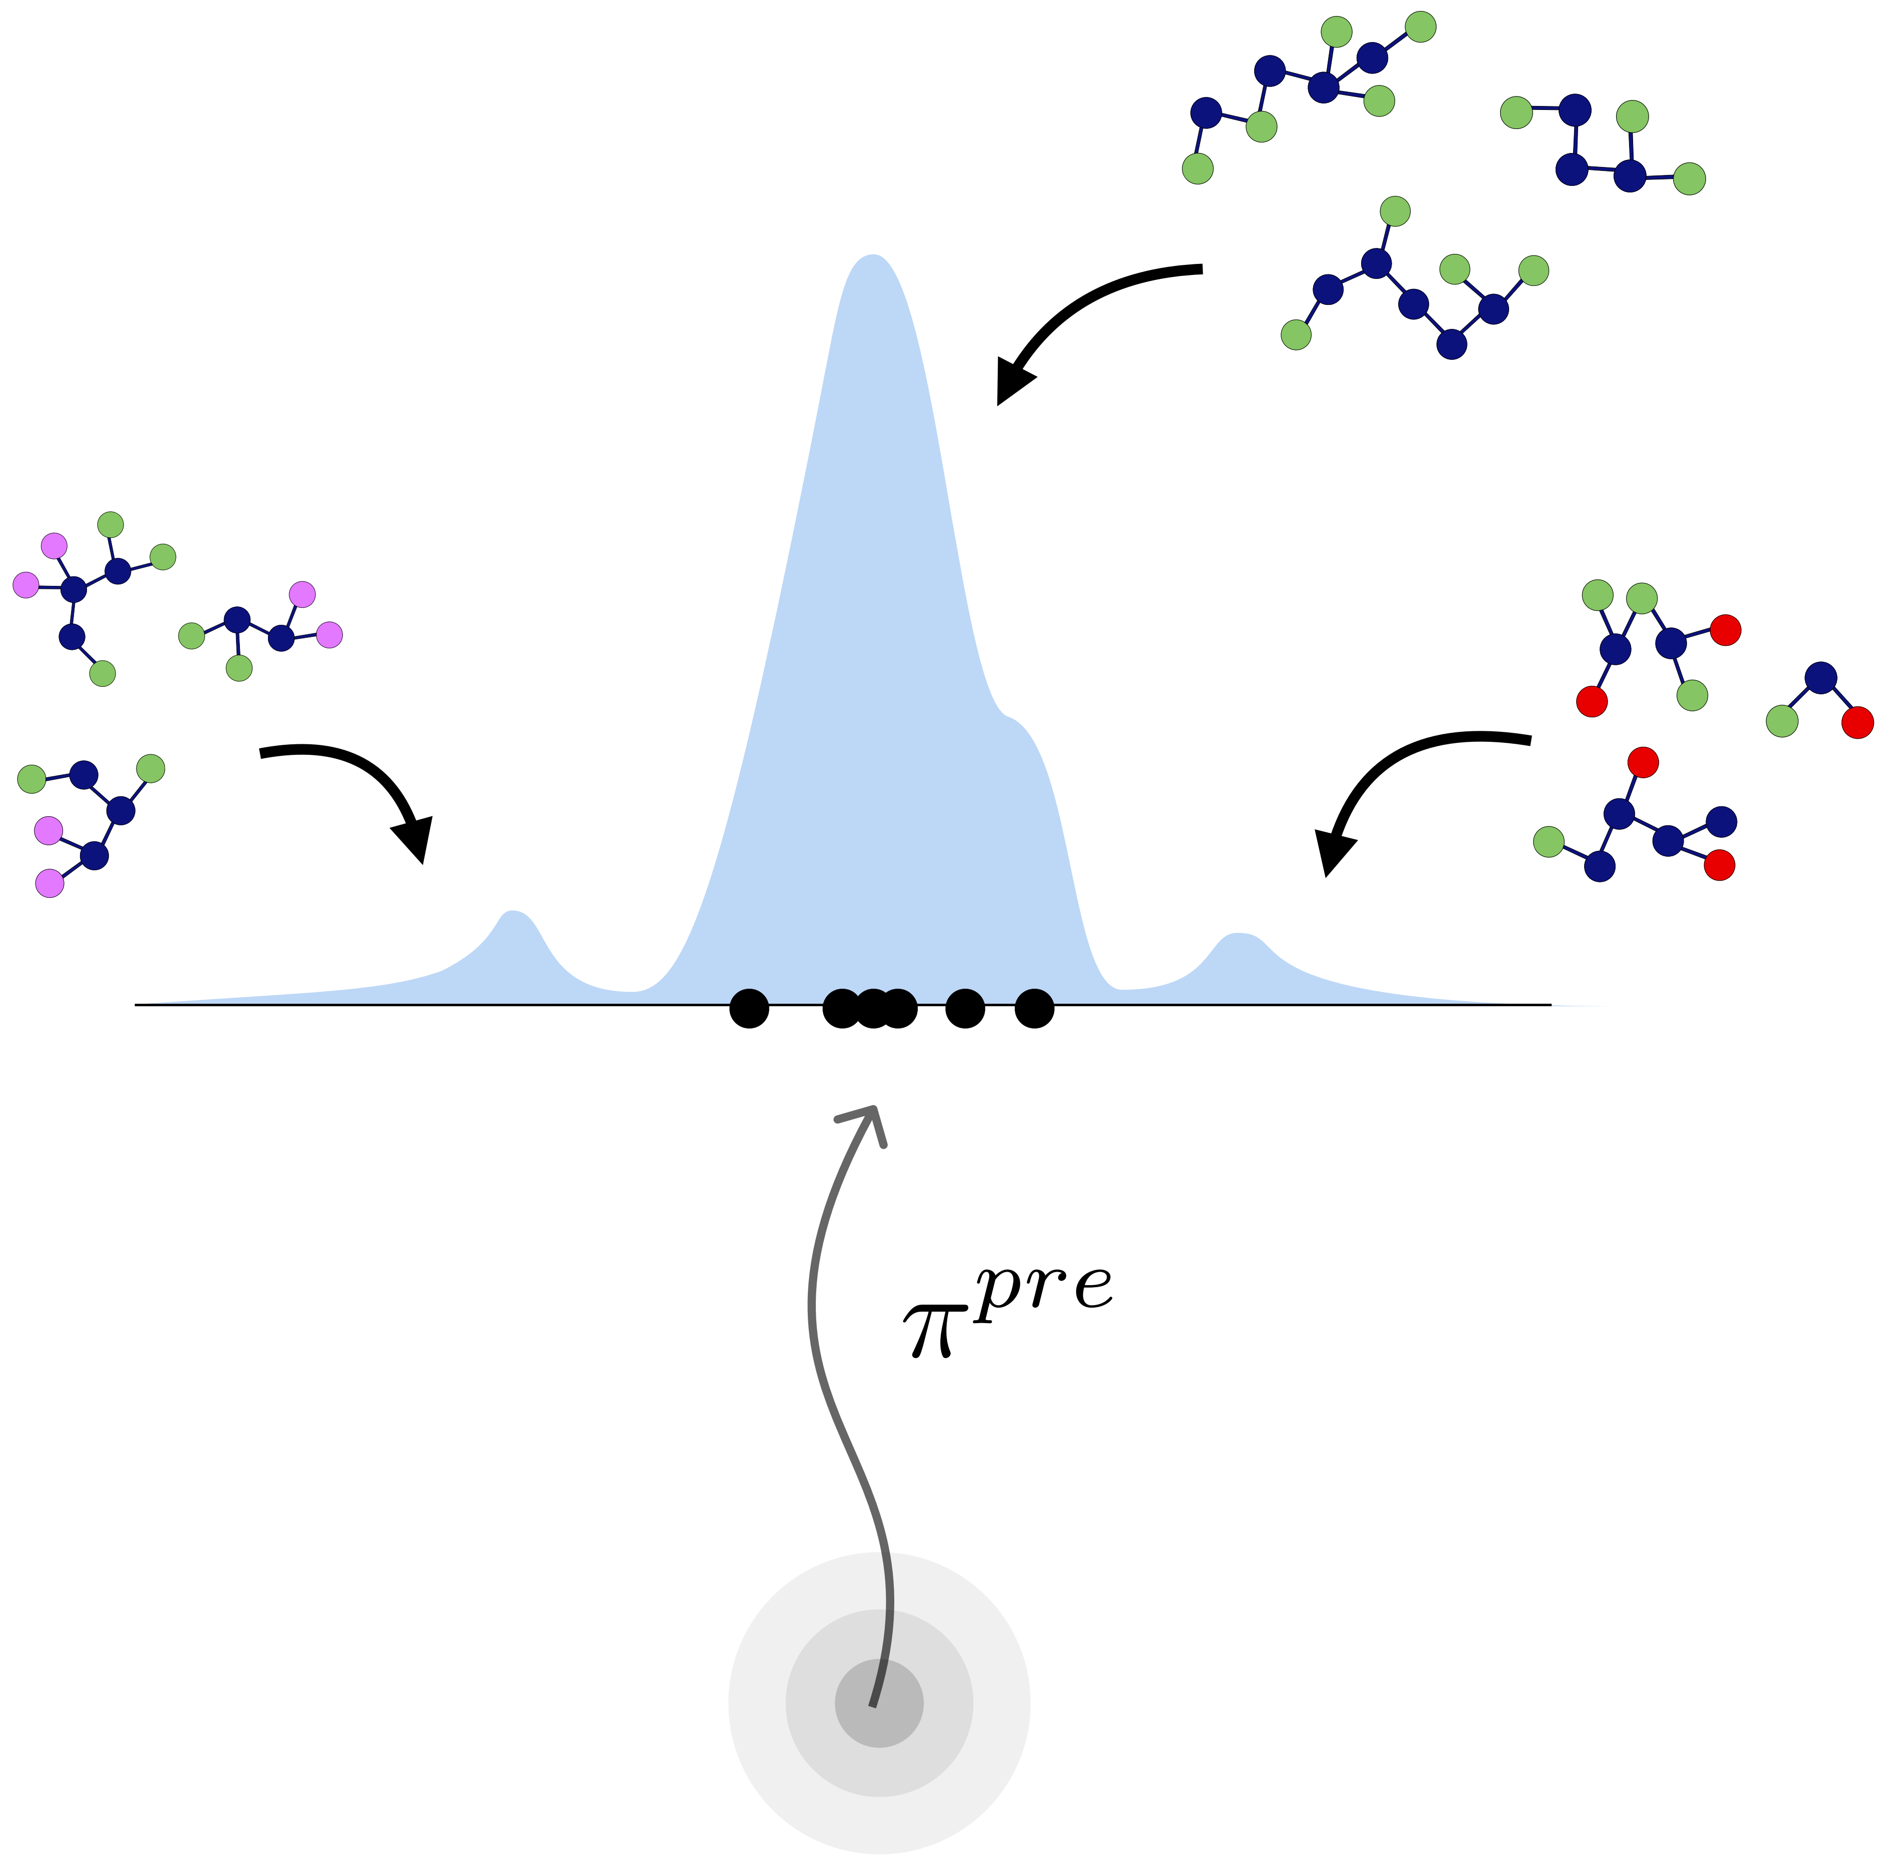

Data Debiasing leads to *Hidden* Mode Discovery

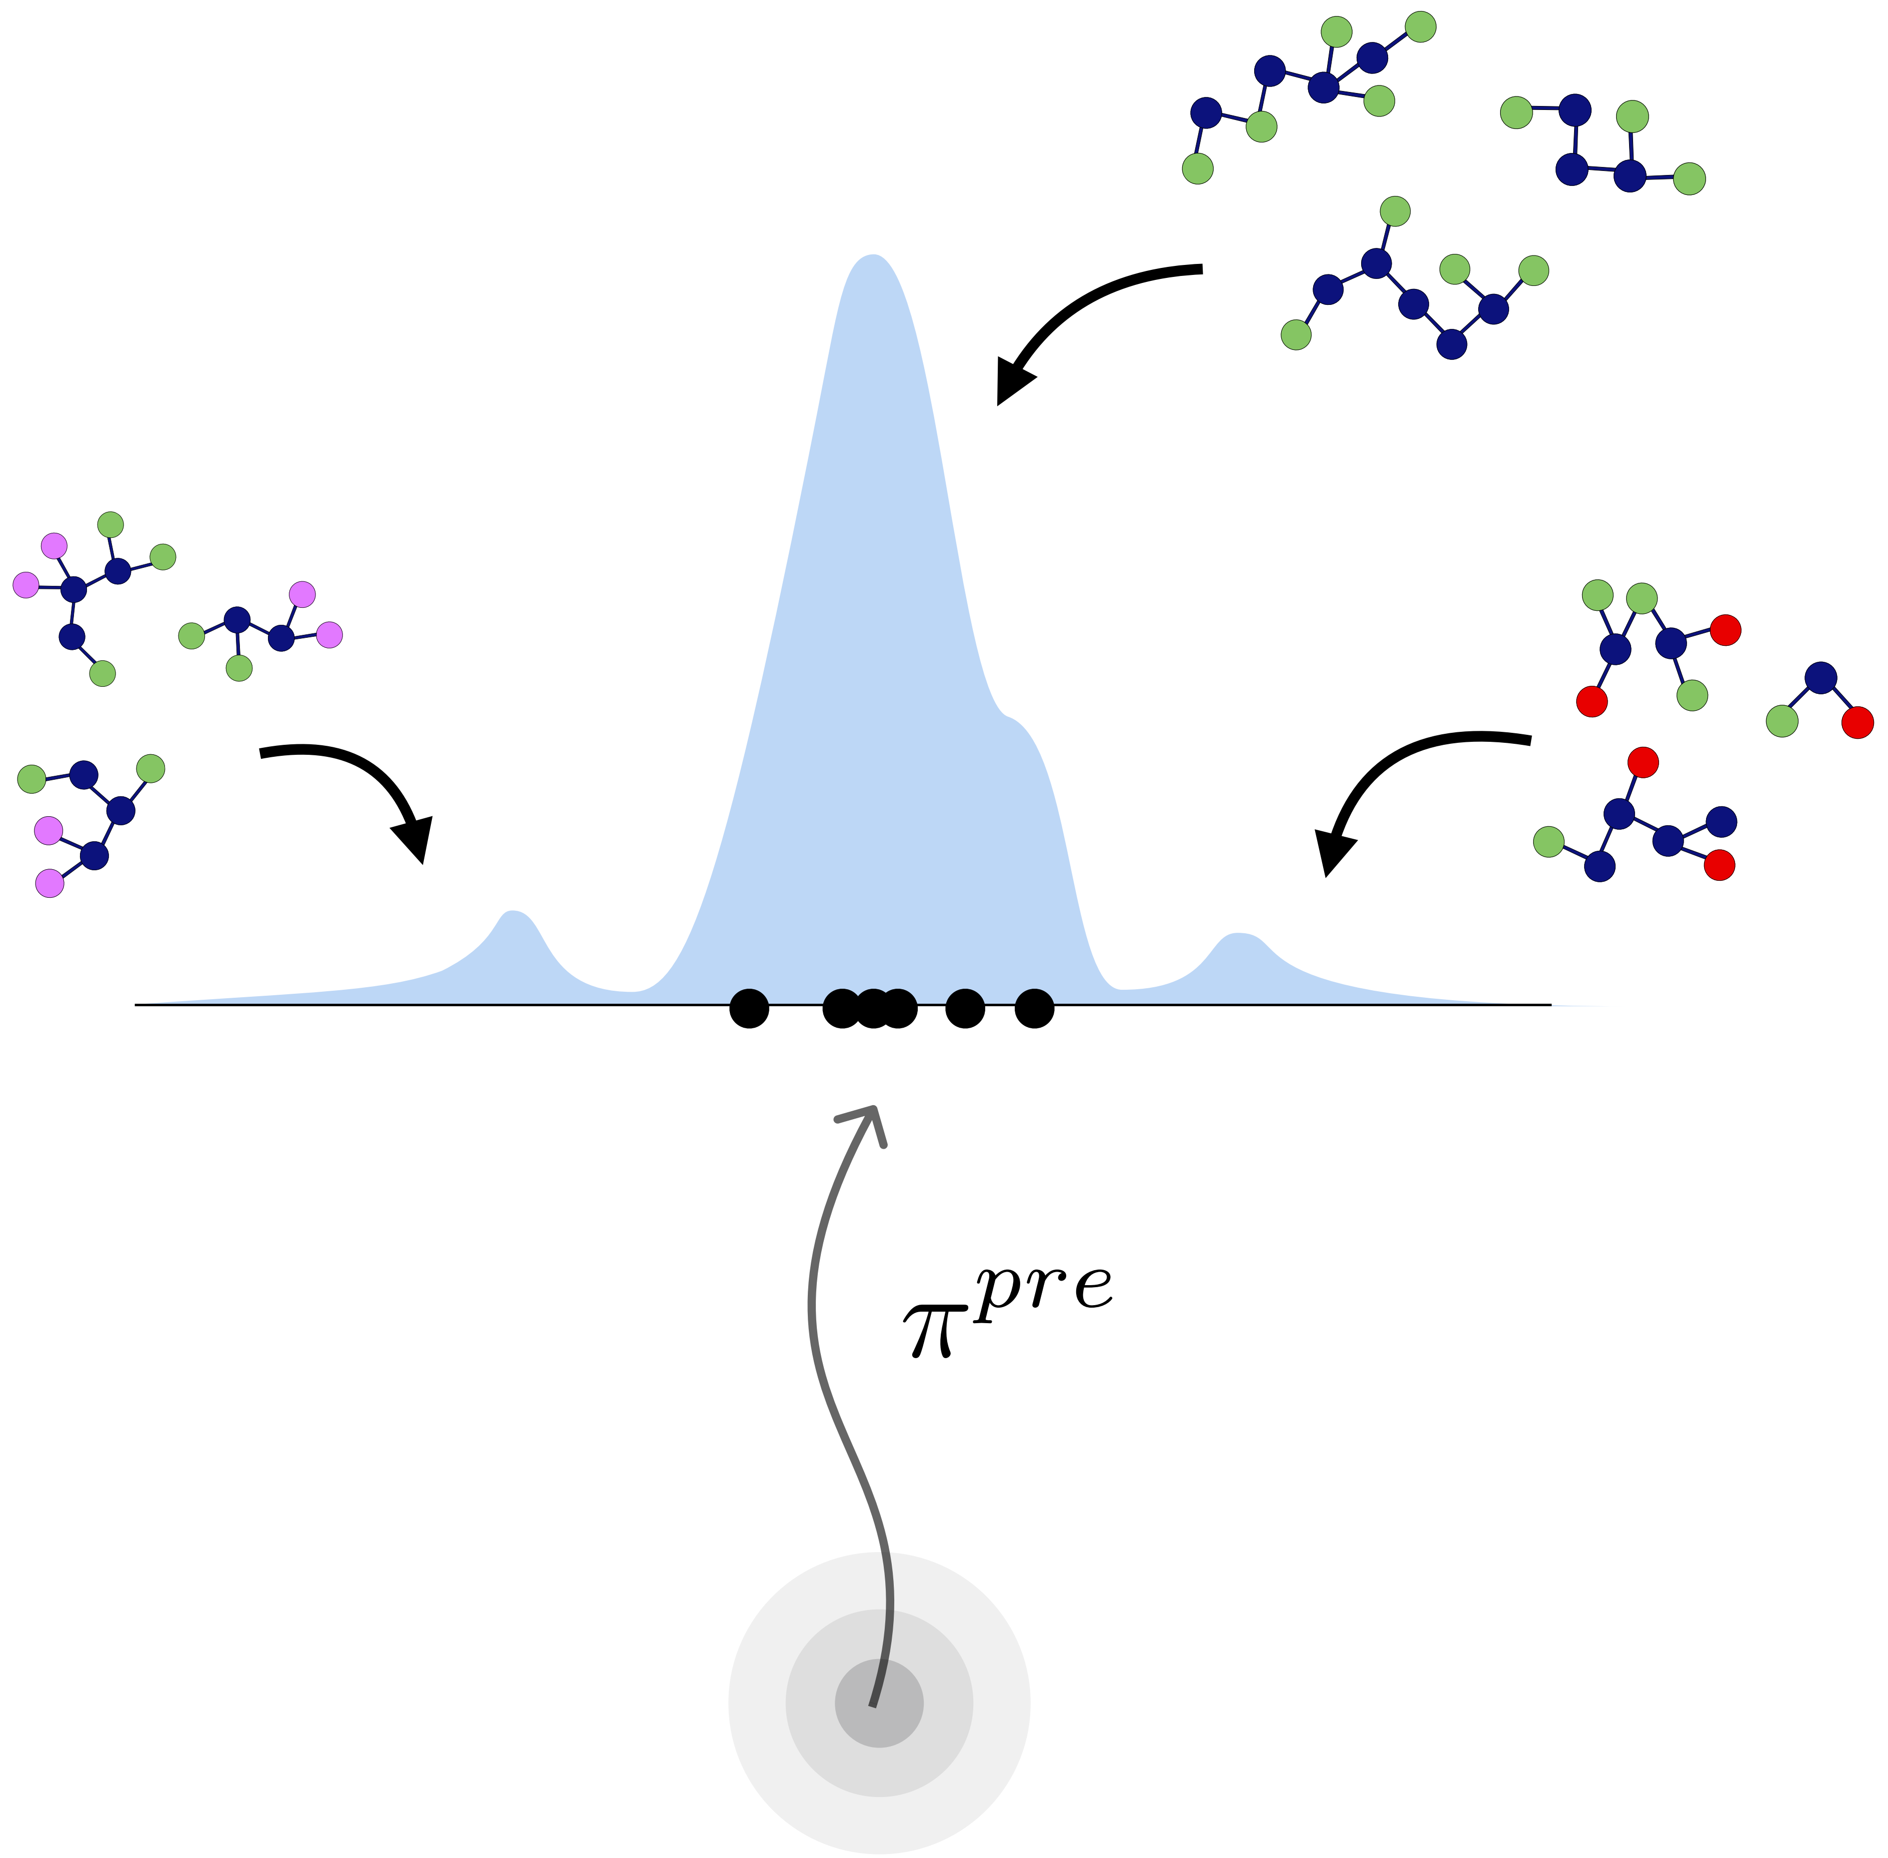

Key Hypothesis: Foundation models learn rich representations, capturing valid low-probability modes.

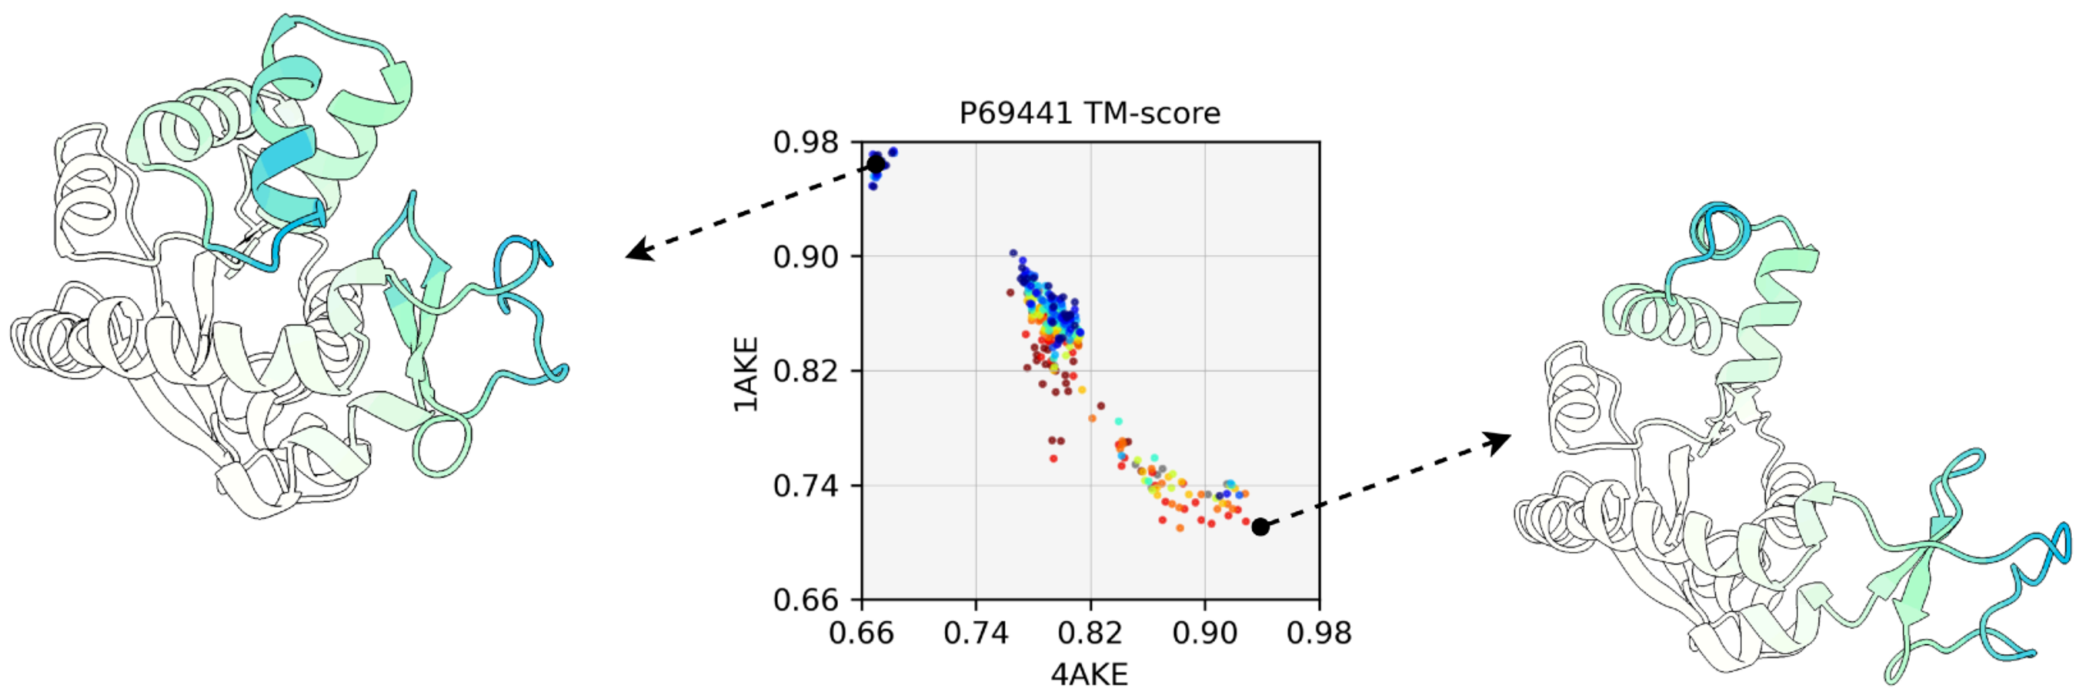

[Shosuke Suzuki et al., 2026]

Data Debiasing leads to *Hidden* Mode Discovery

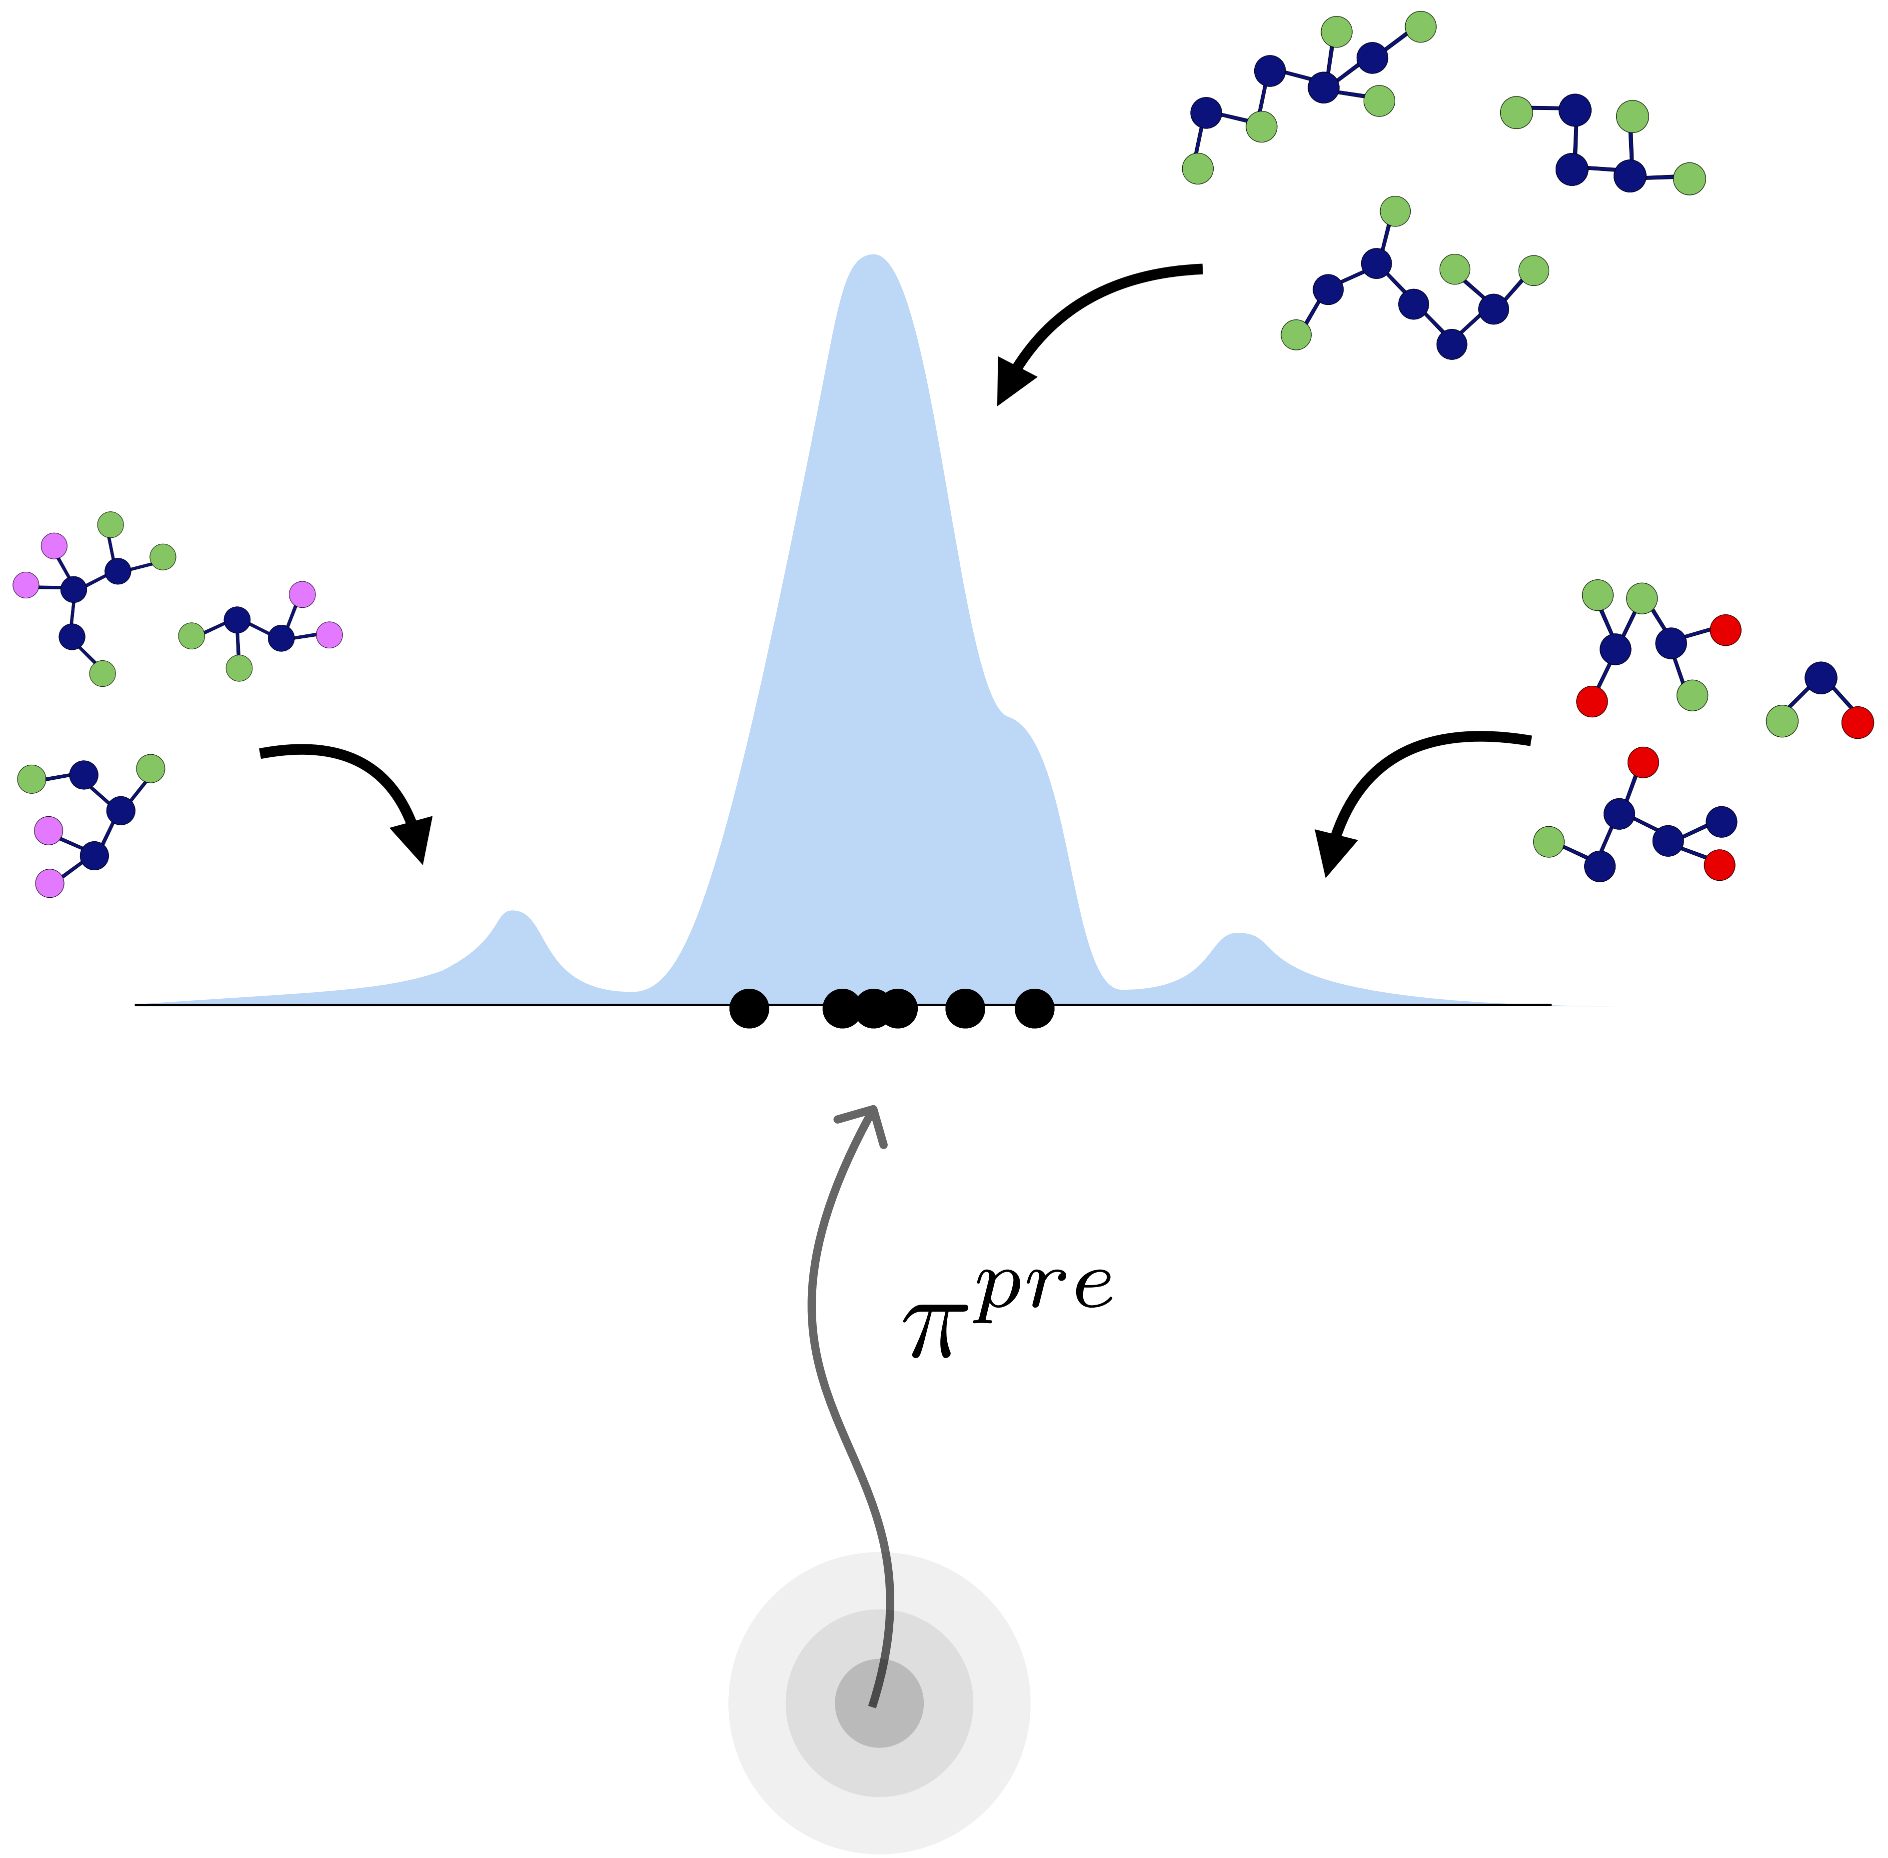

Key Hypothesis: Foundation models learn rich representations, capturing valid low-probability modes.

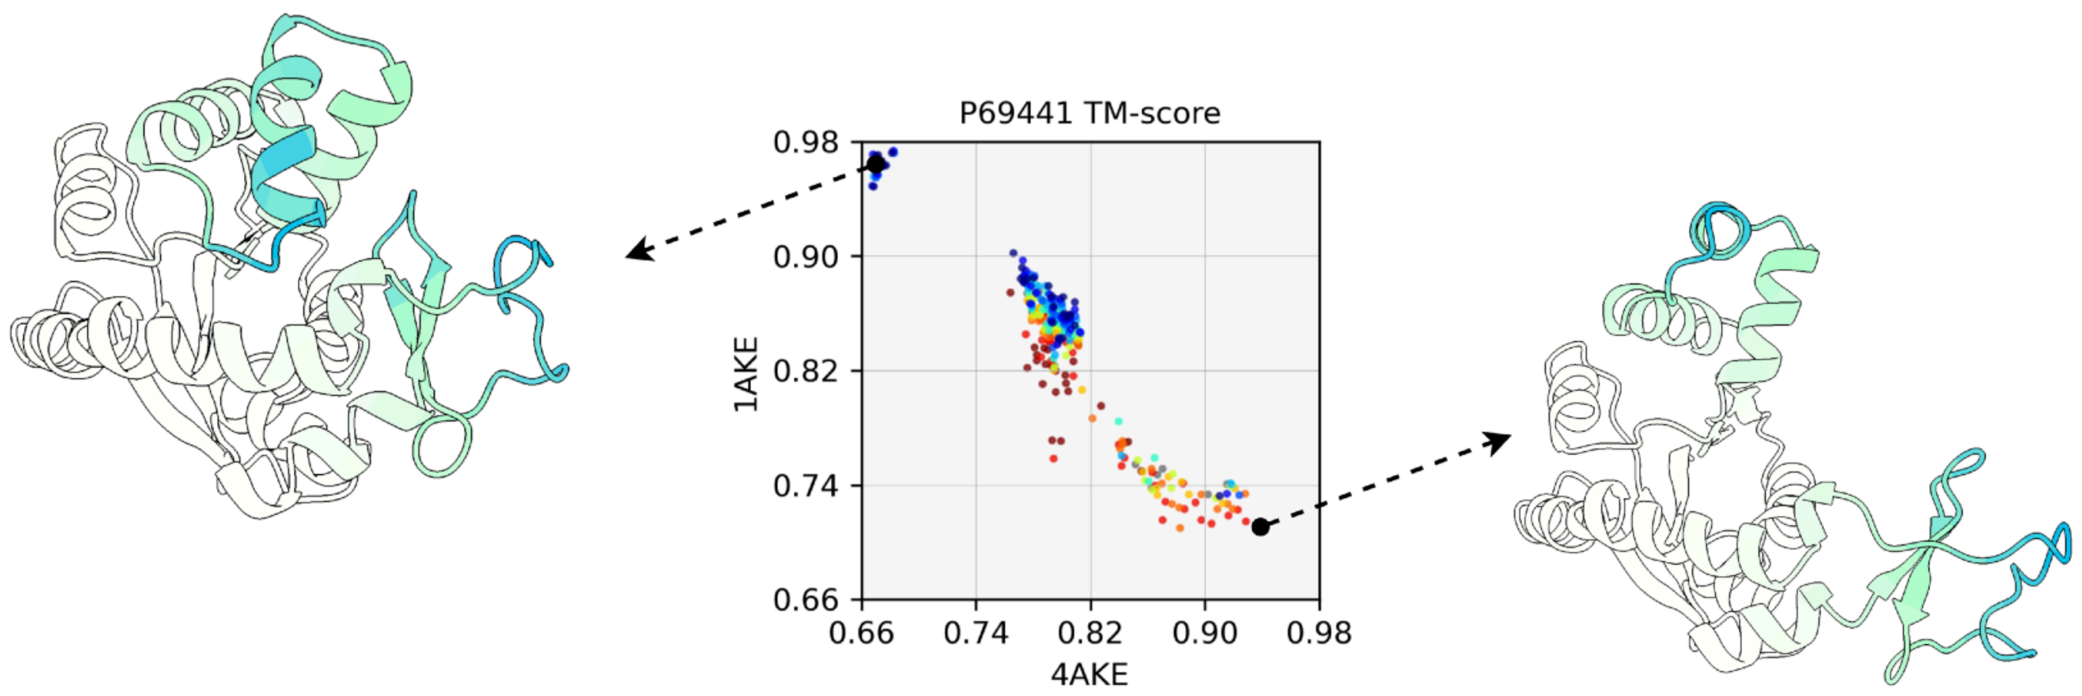

Finite-budget sampling rarely leads to samples from low-probability modes!

⟶ **hidden modes** when generating samples

[Shosuke Suzuki et al., 2026]

Flow Density Control for Entropic Expansion

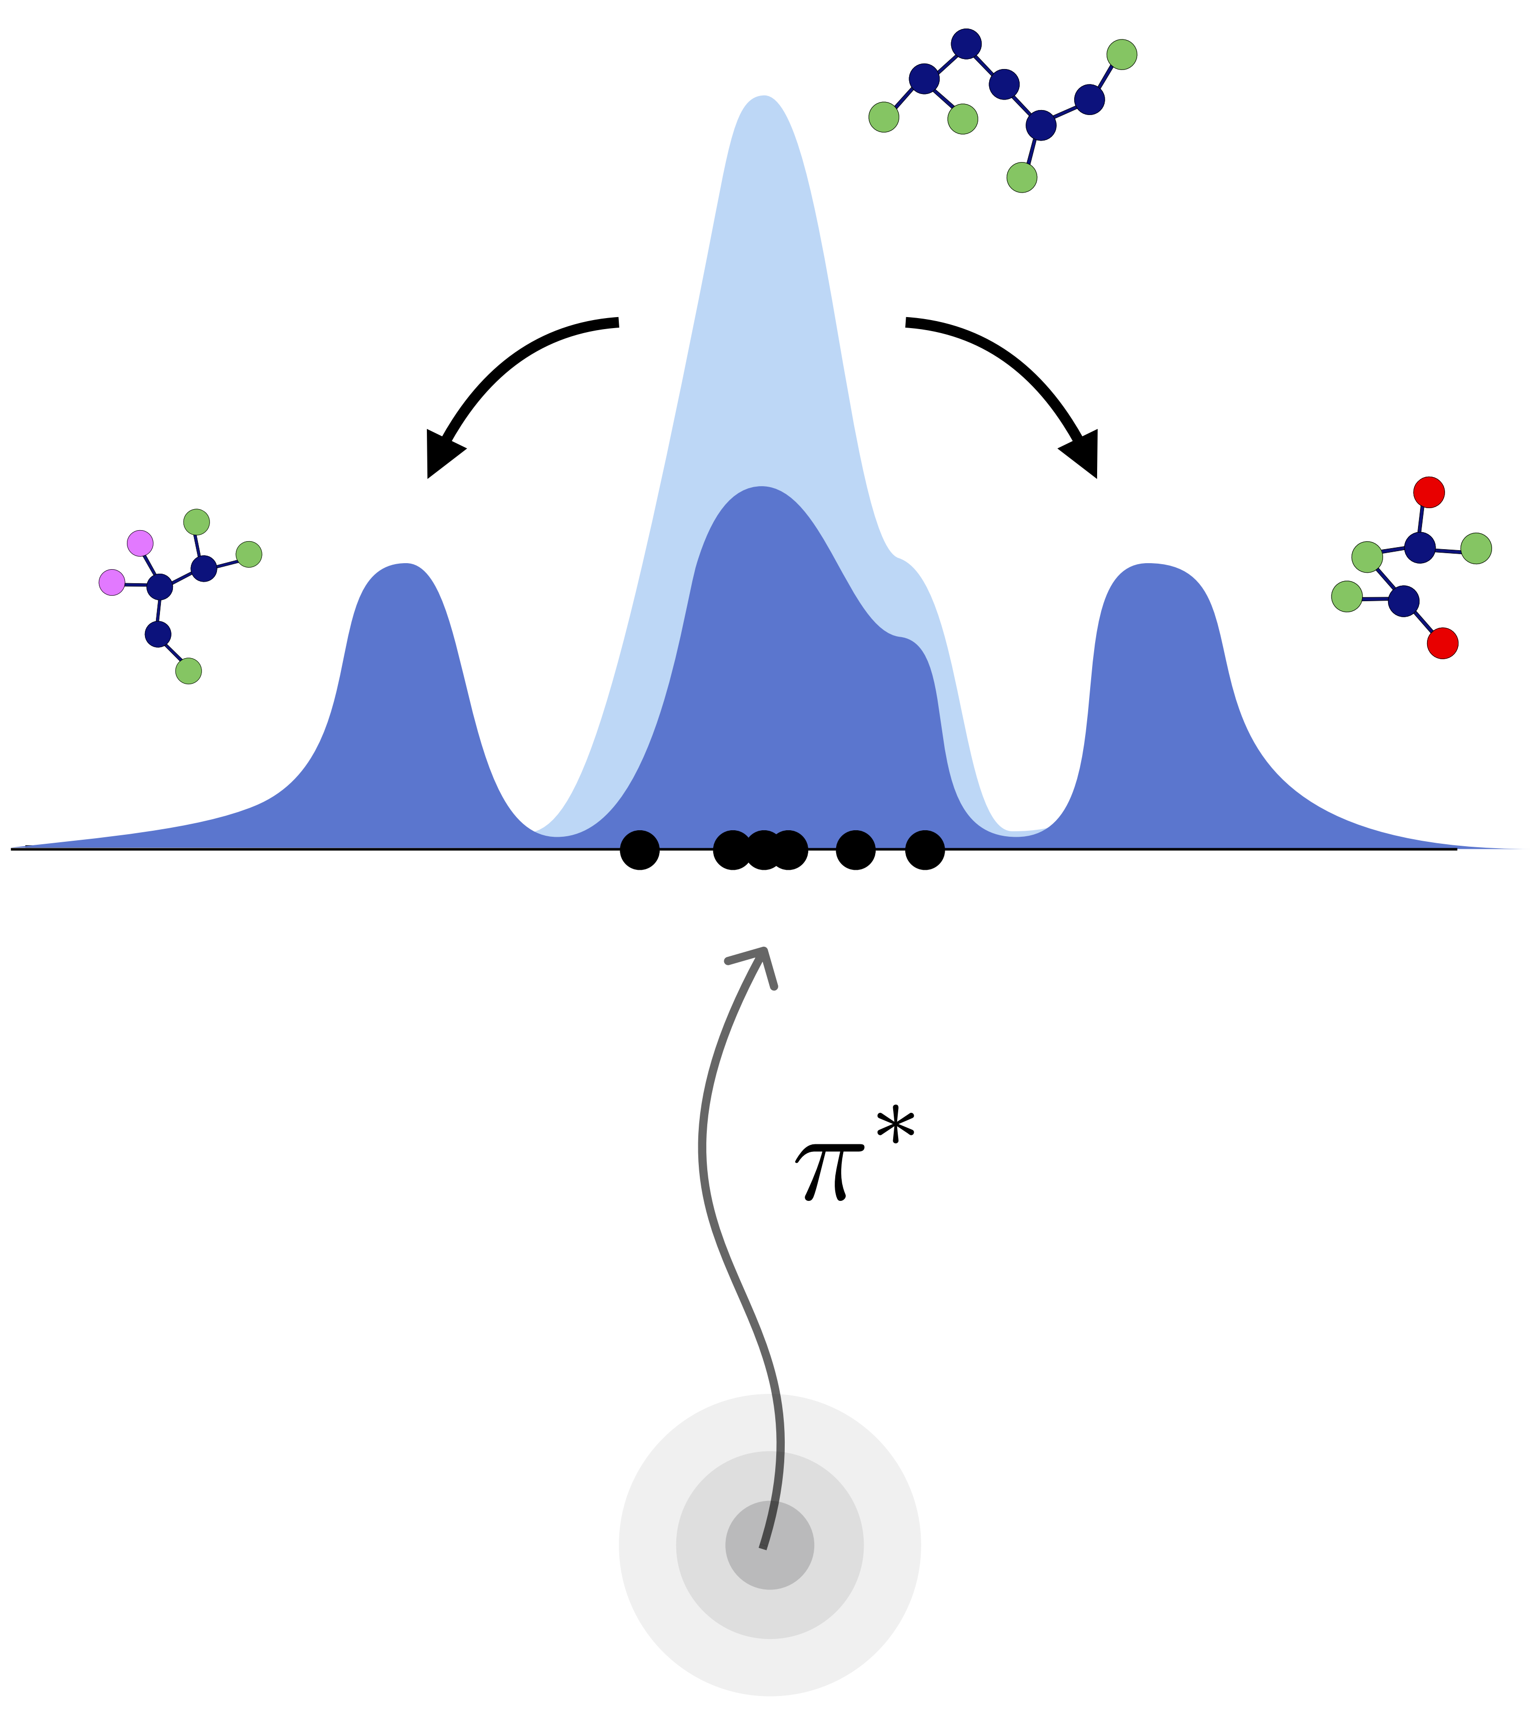

KL-regularized Entropic Expansion

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

$\mathcal{H}(p_1^\pi)$ → **Increases diversity**

$\alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre})$ → **Avoid invalid regions**

Flow Density Control for Entropic Expansion

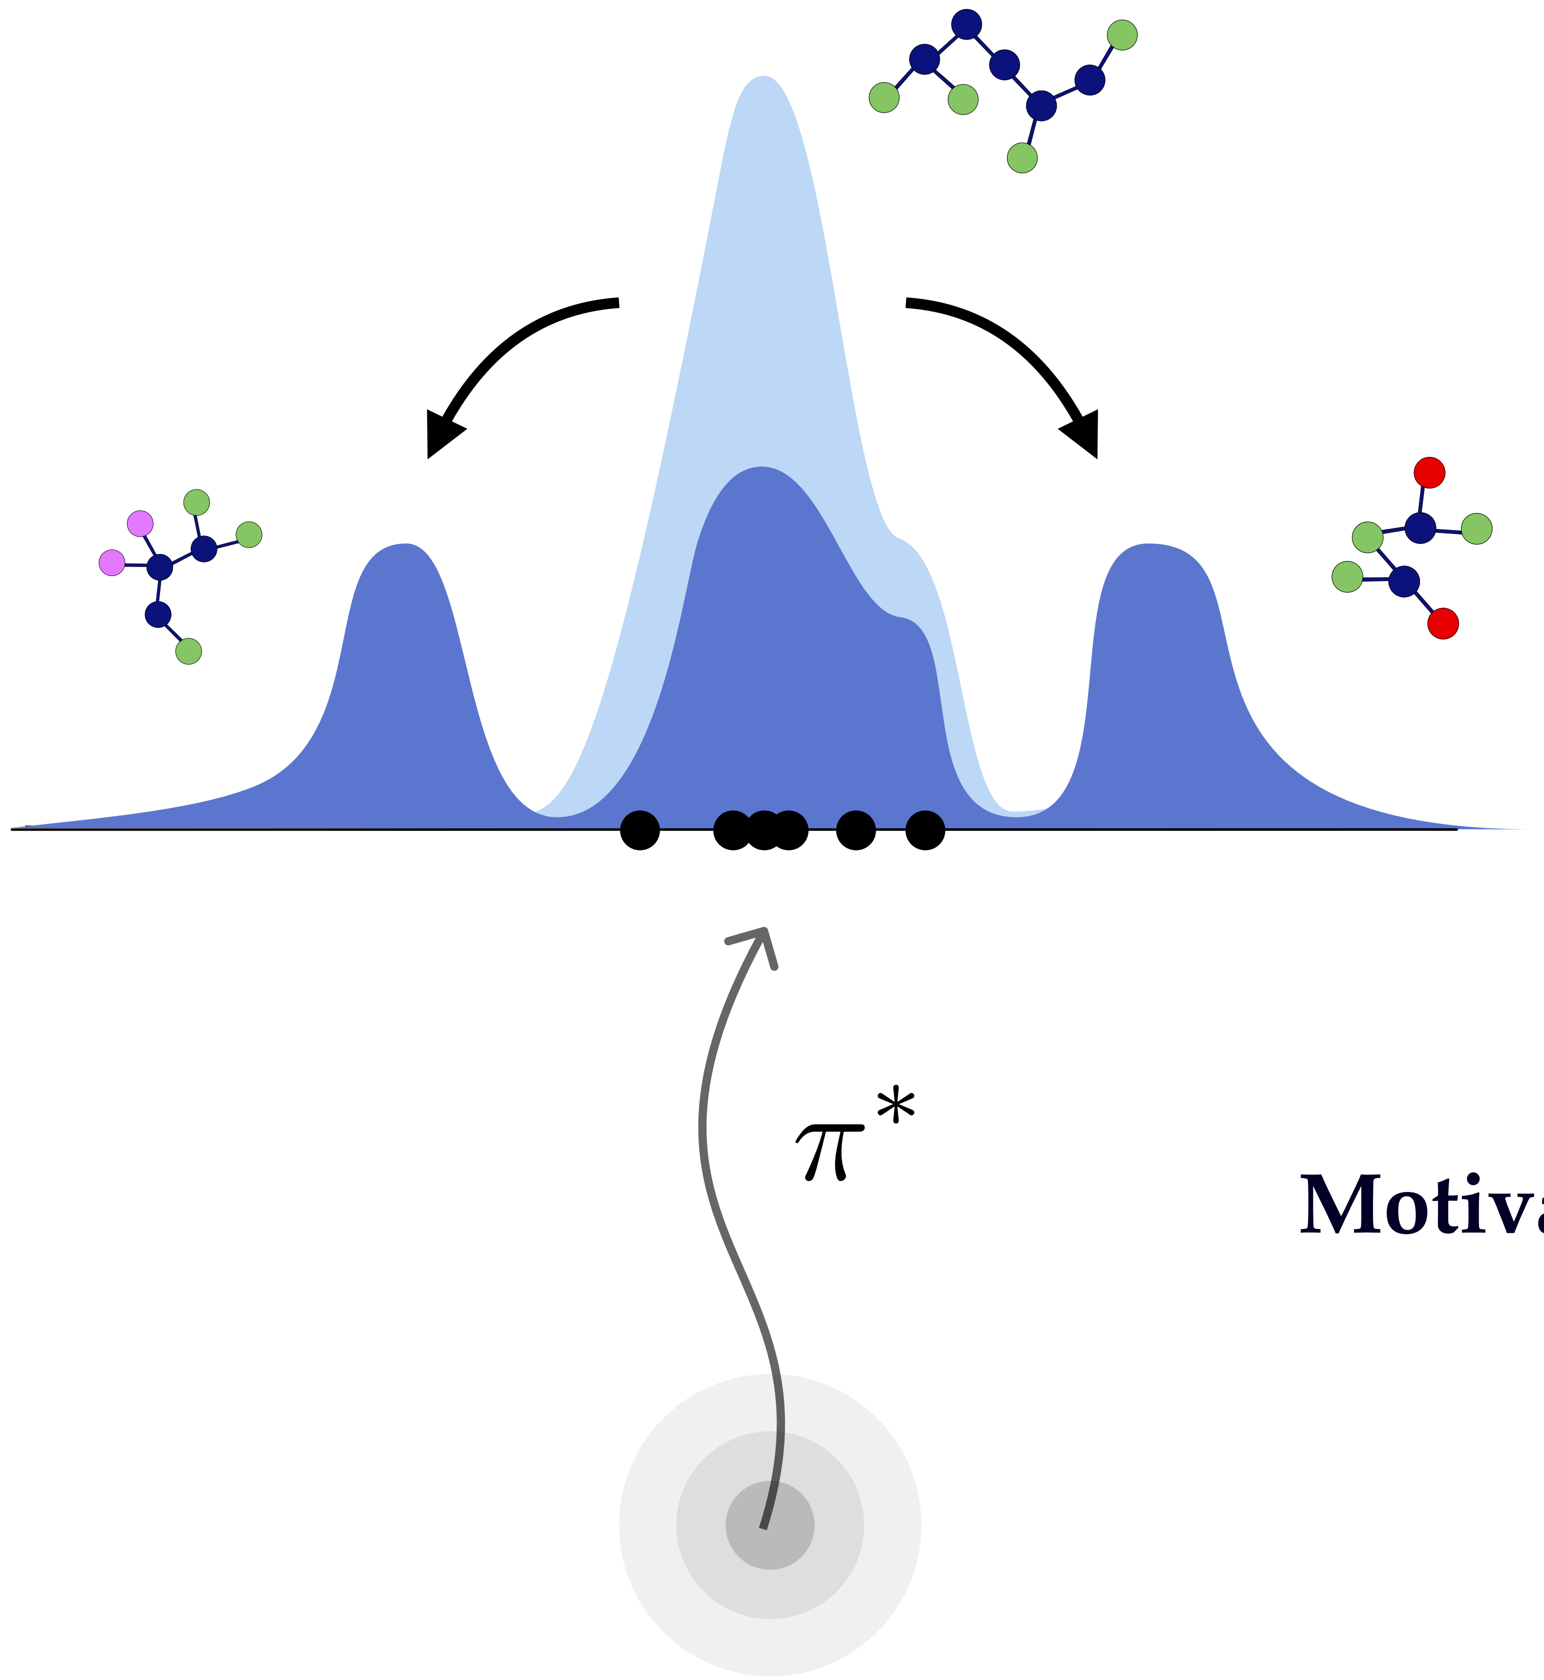

KL-regularized Entropic Expansion

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

Increases diversity

Avoid invalid regions

Motivation:

- **entropy** allows to **re-distribute density over the model's *hidden* modes**
- **KL-regularization** prevents transport to likely invalid regions

→ **access novel regions** where the model places insufficient mass

How to Increase Flow Entropy via RL

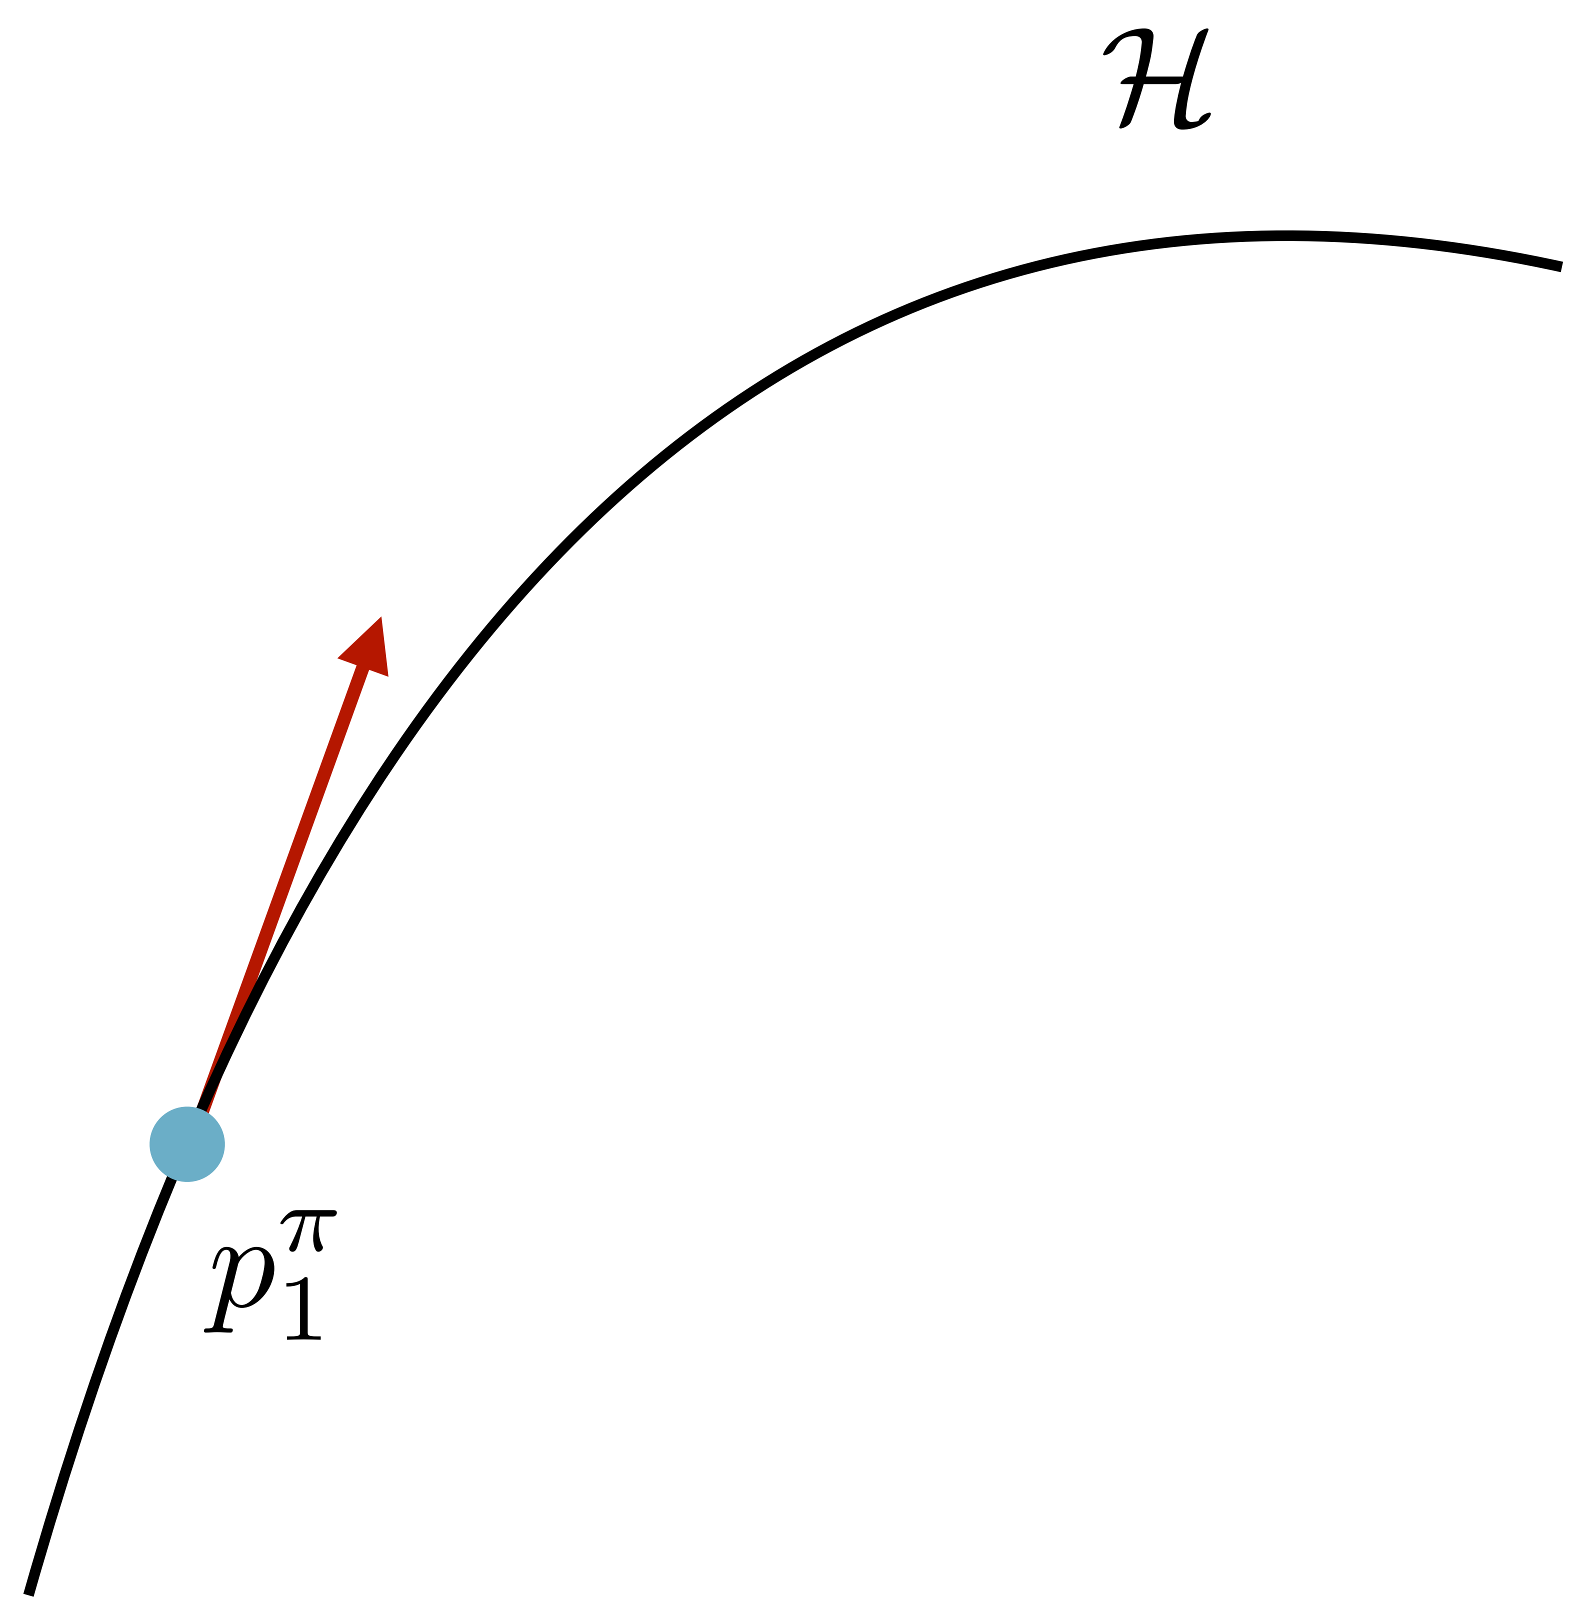

How to Increase Flow Entropy via RL

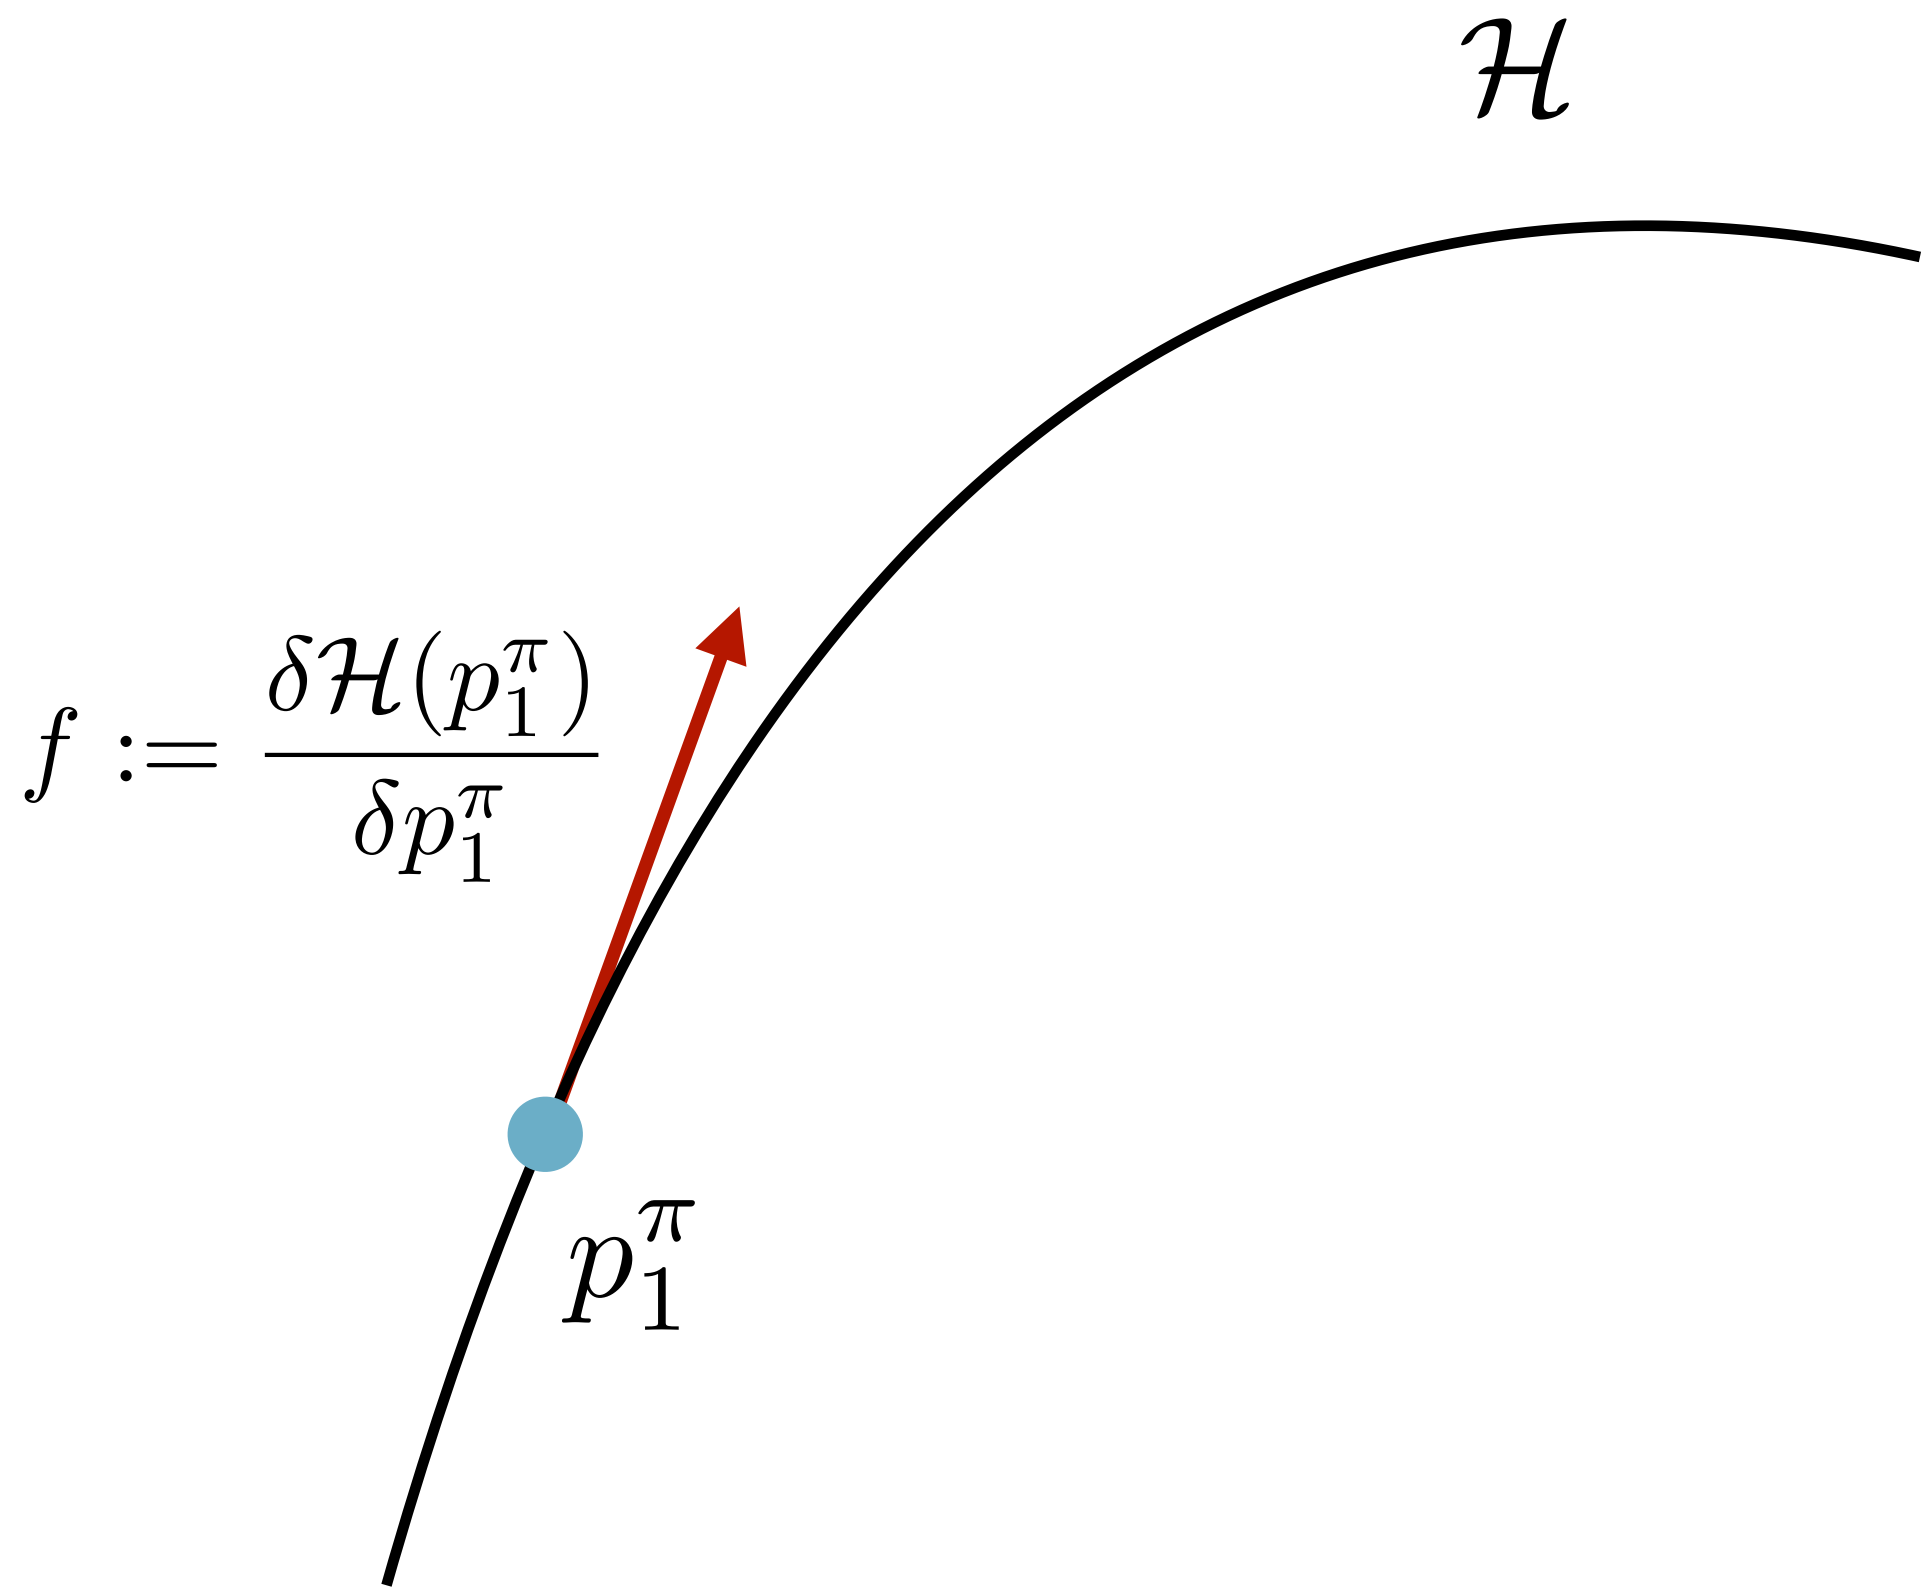

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

How to Increase Flow Entropy via RL

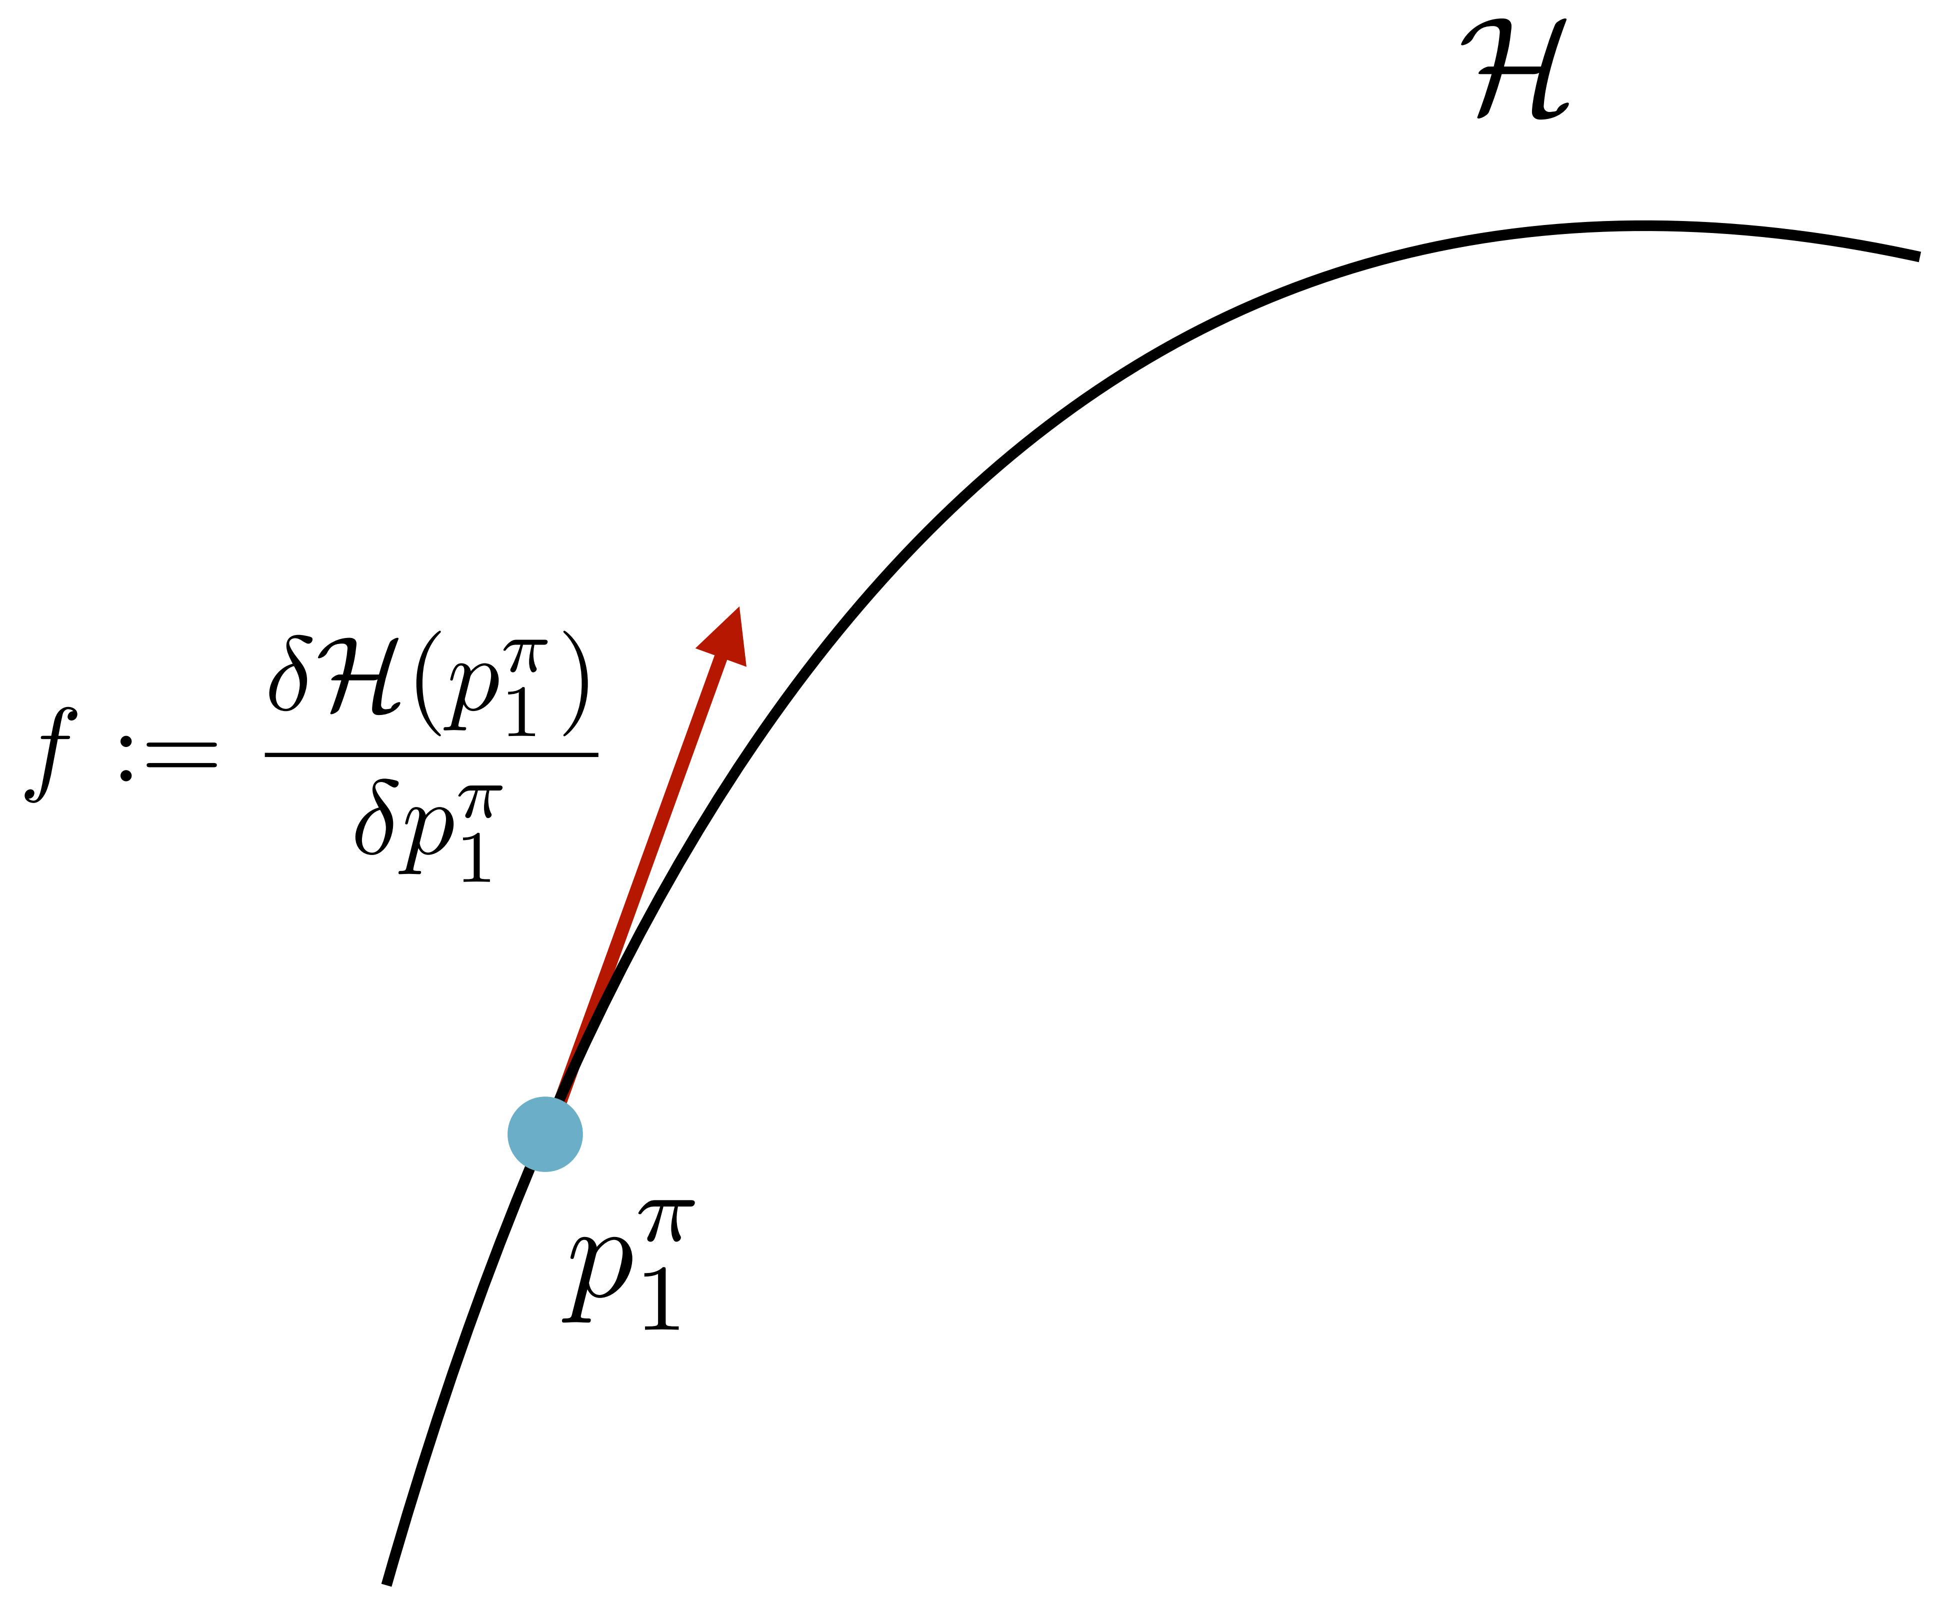

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Log likelihoods are hard to estimate!

(currently does not scale)

How to Increase Flow Entropy via RL

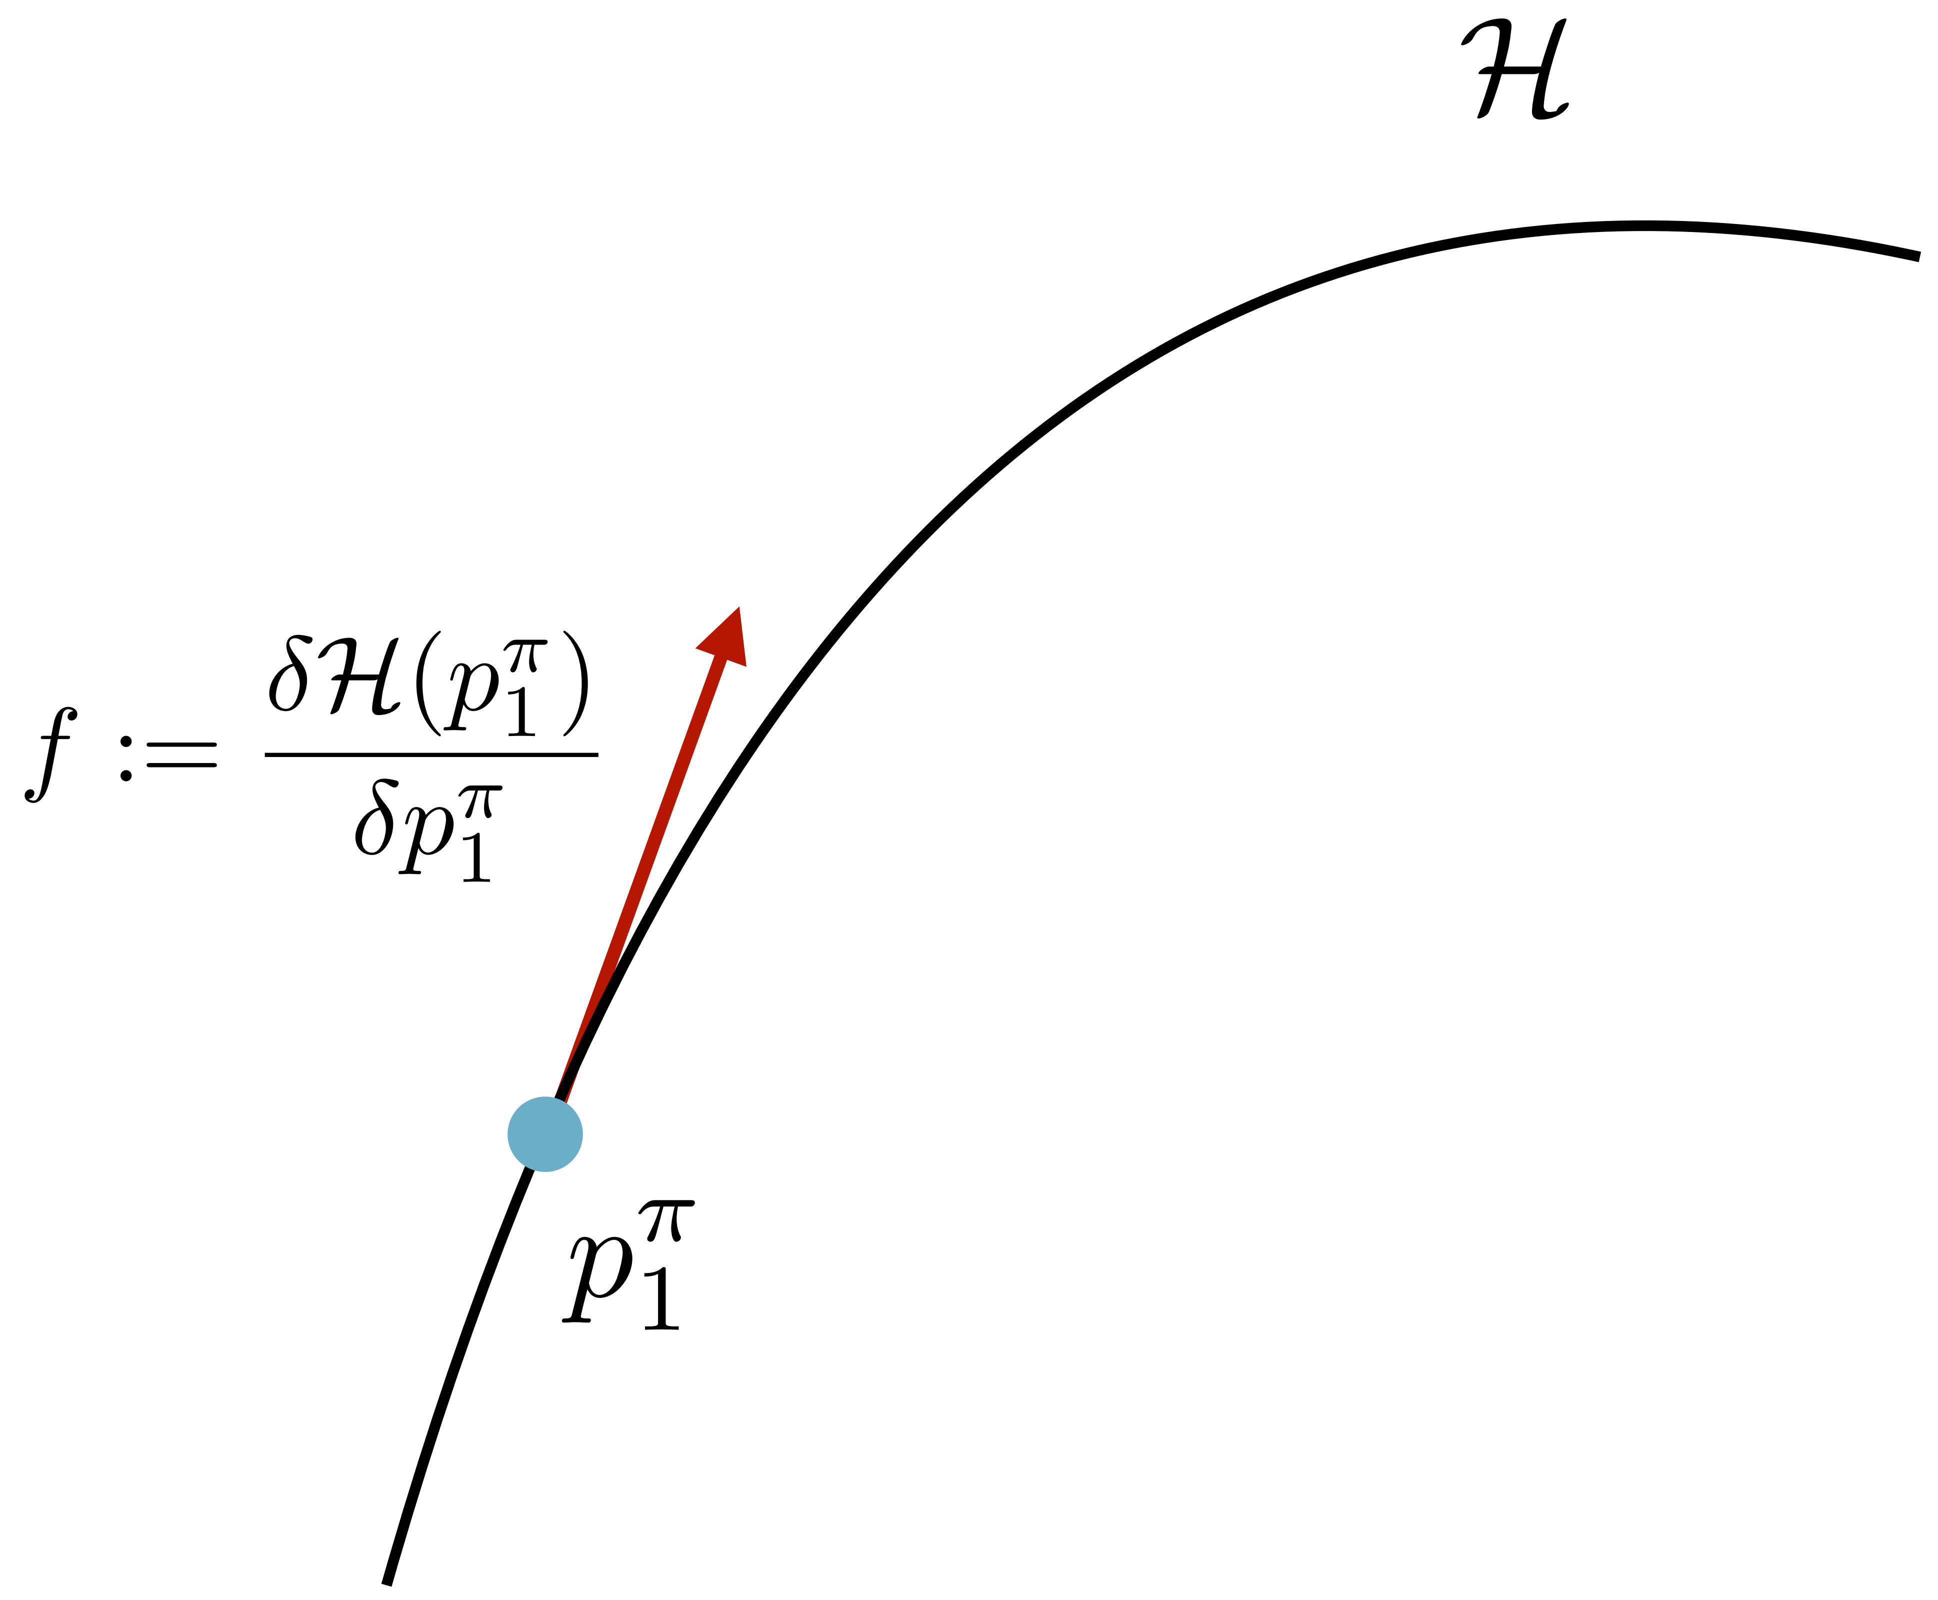

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Log likelihoods are hard to estimate!

(currently does not scale)

We can fully bypass log-likelihood estimation!

How to Increase Flow Entropy via RL

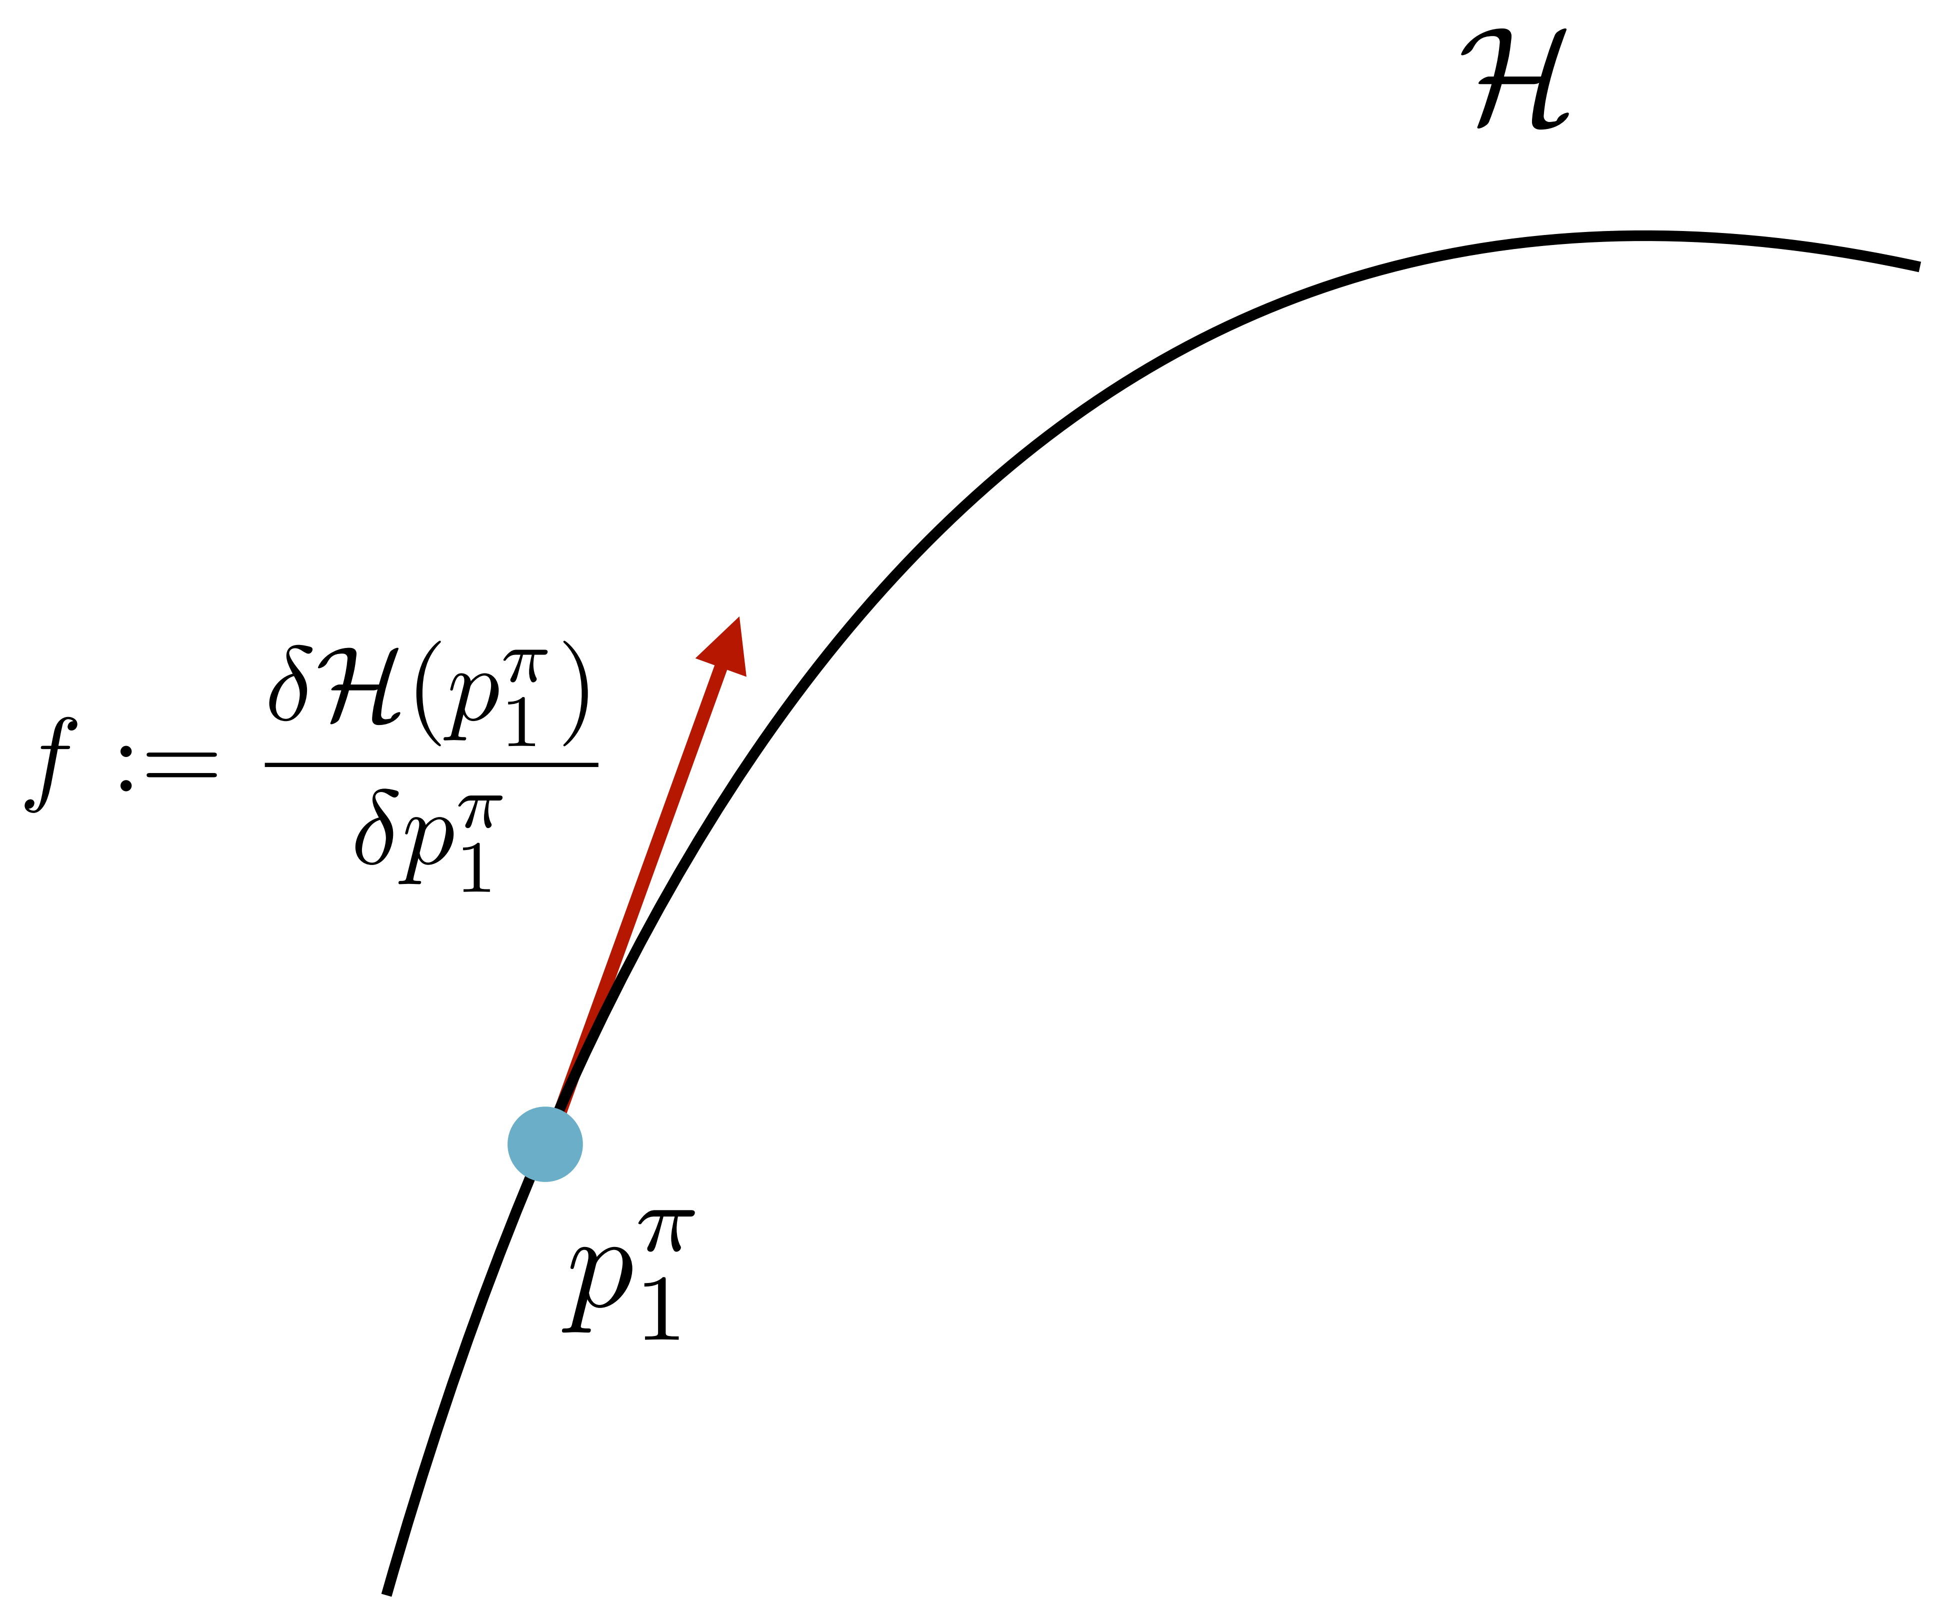

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Closed-form Gradient of Surprise

$$\nabla f = -\nabla \log\left(p_1^\pi\right)$$

How to Increase Flow Entropy via RL

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

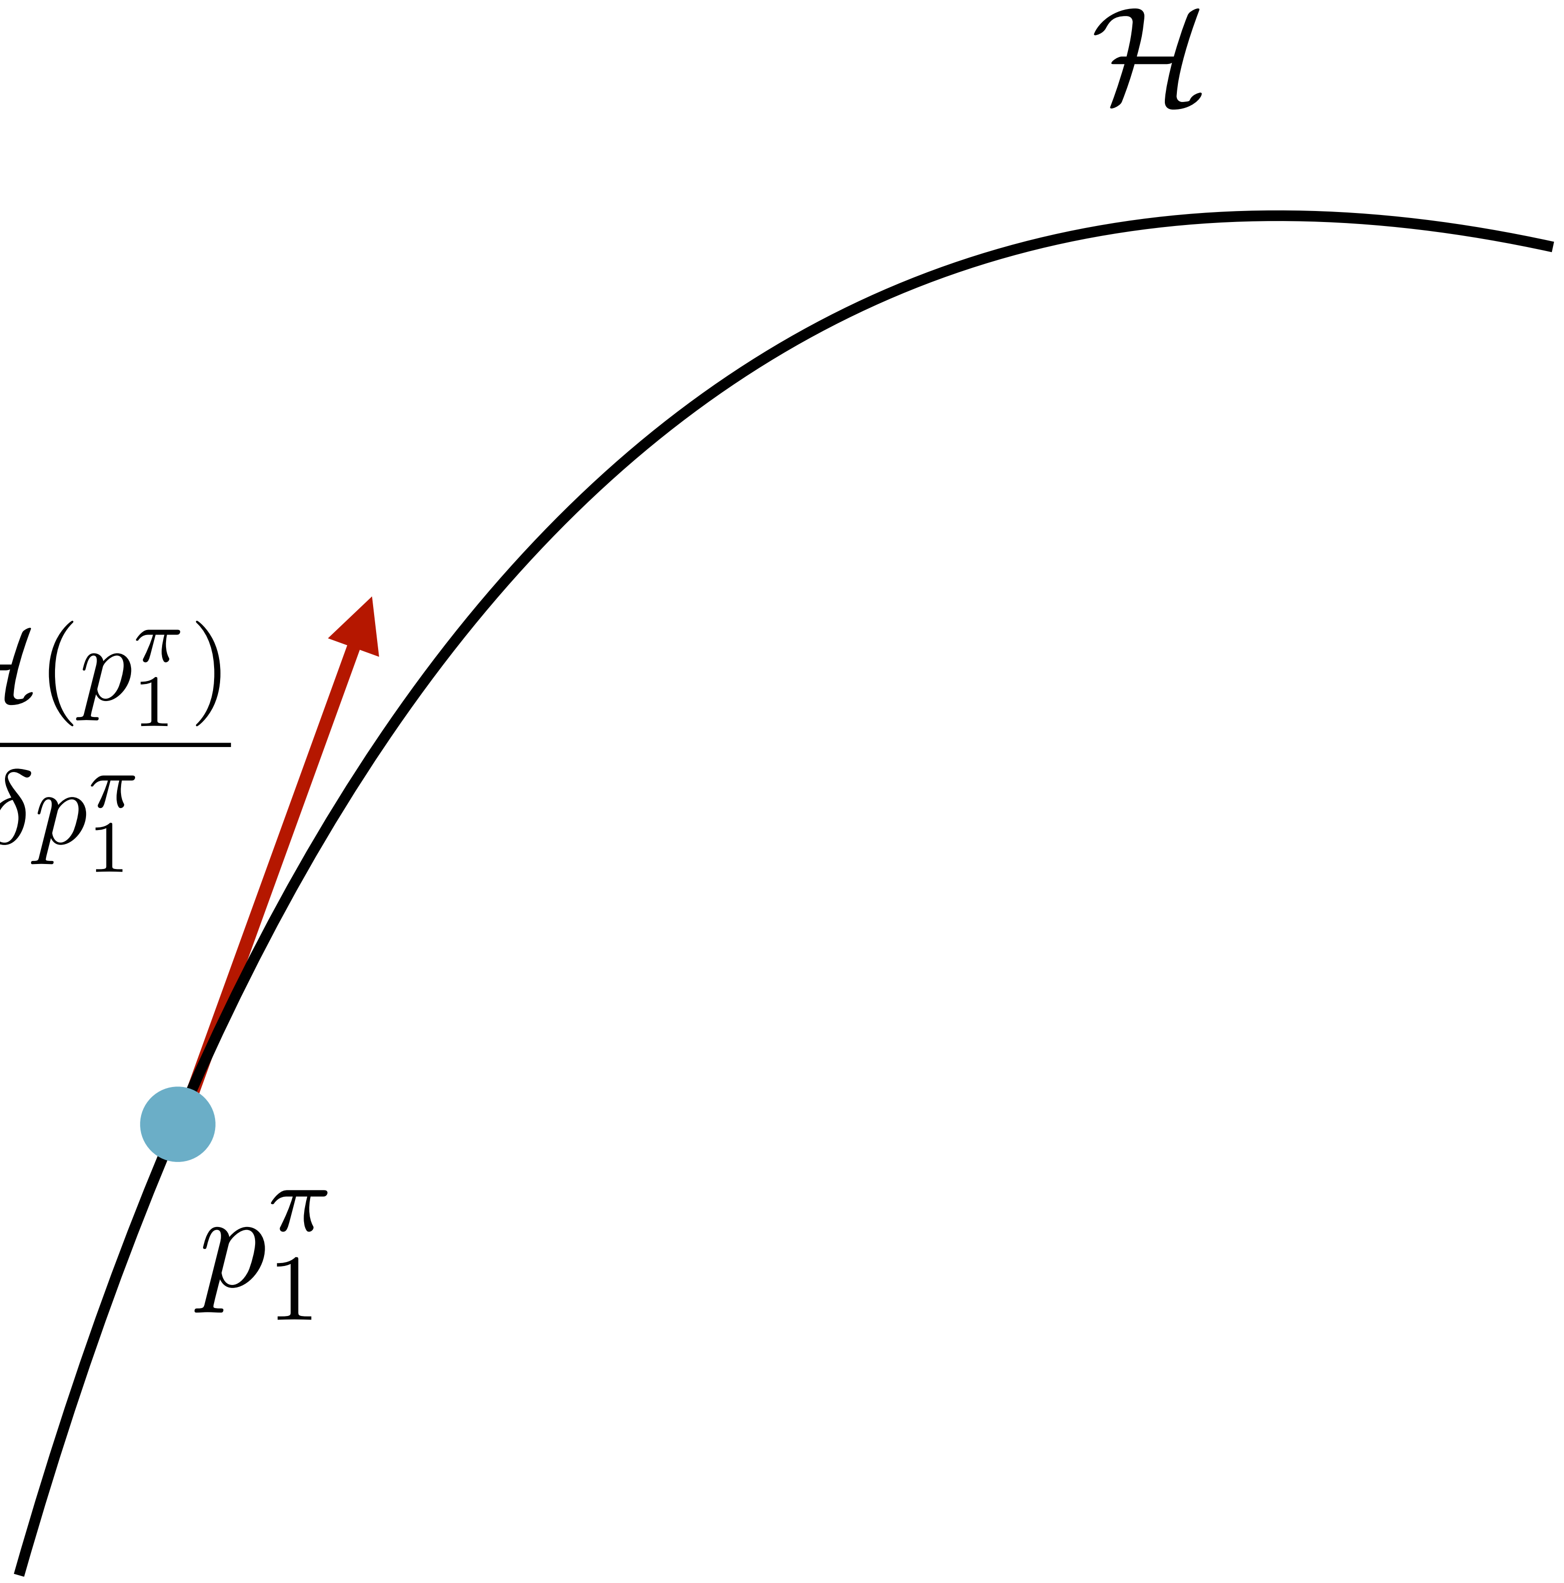

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Closed-form Gradient of Surprise

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

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

How to Increase Flow Entropy via RL

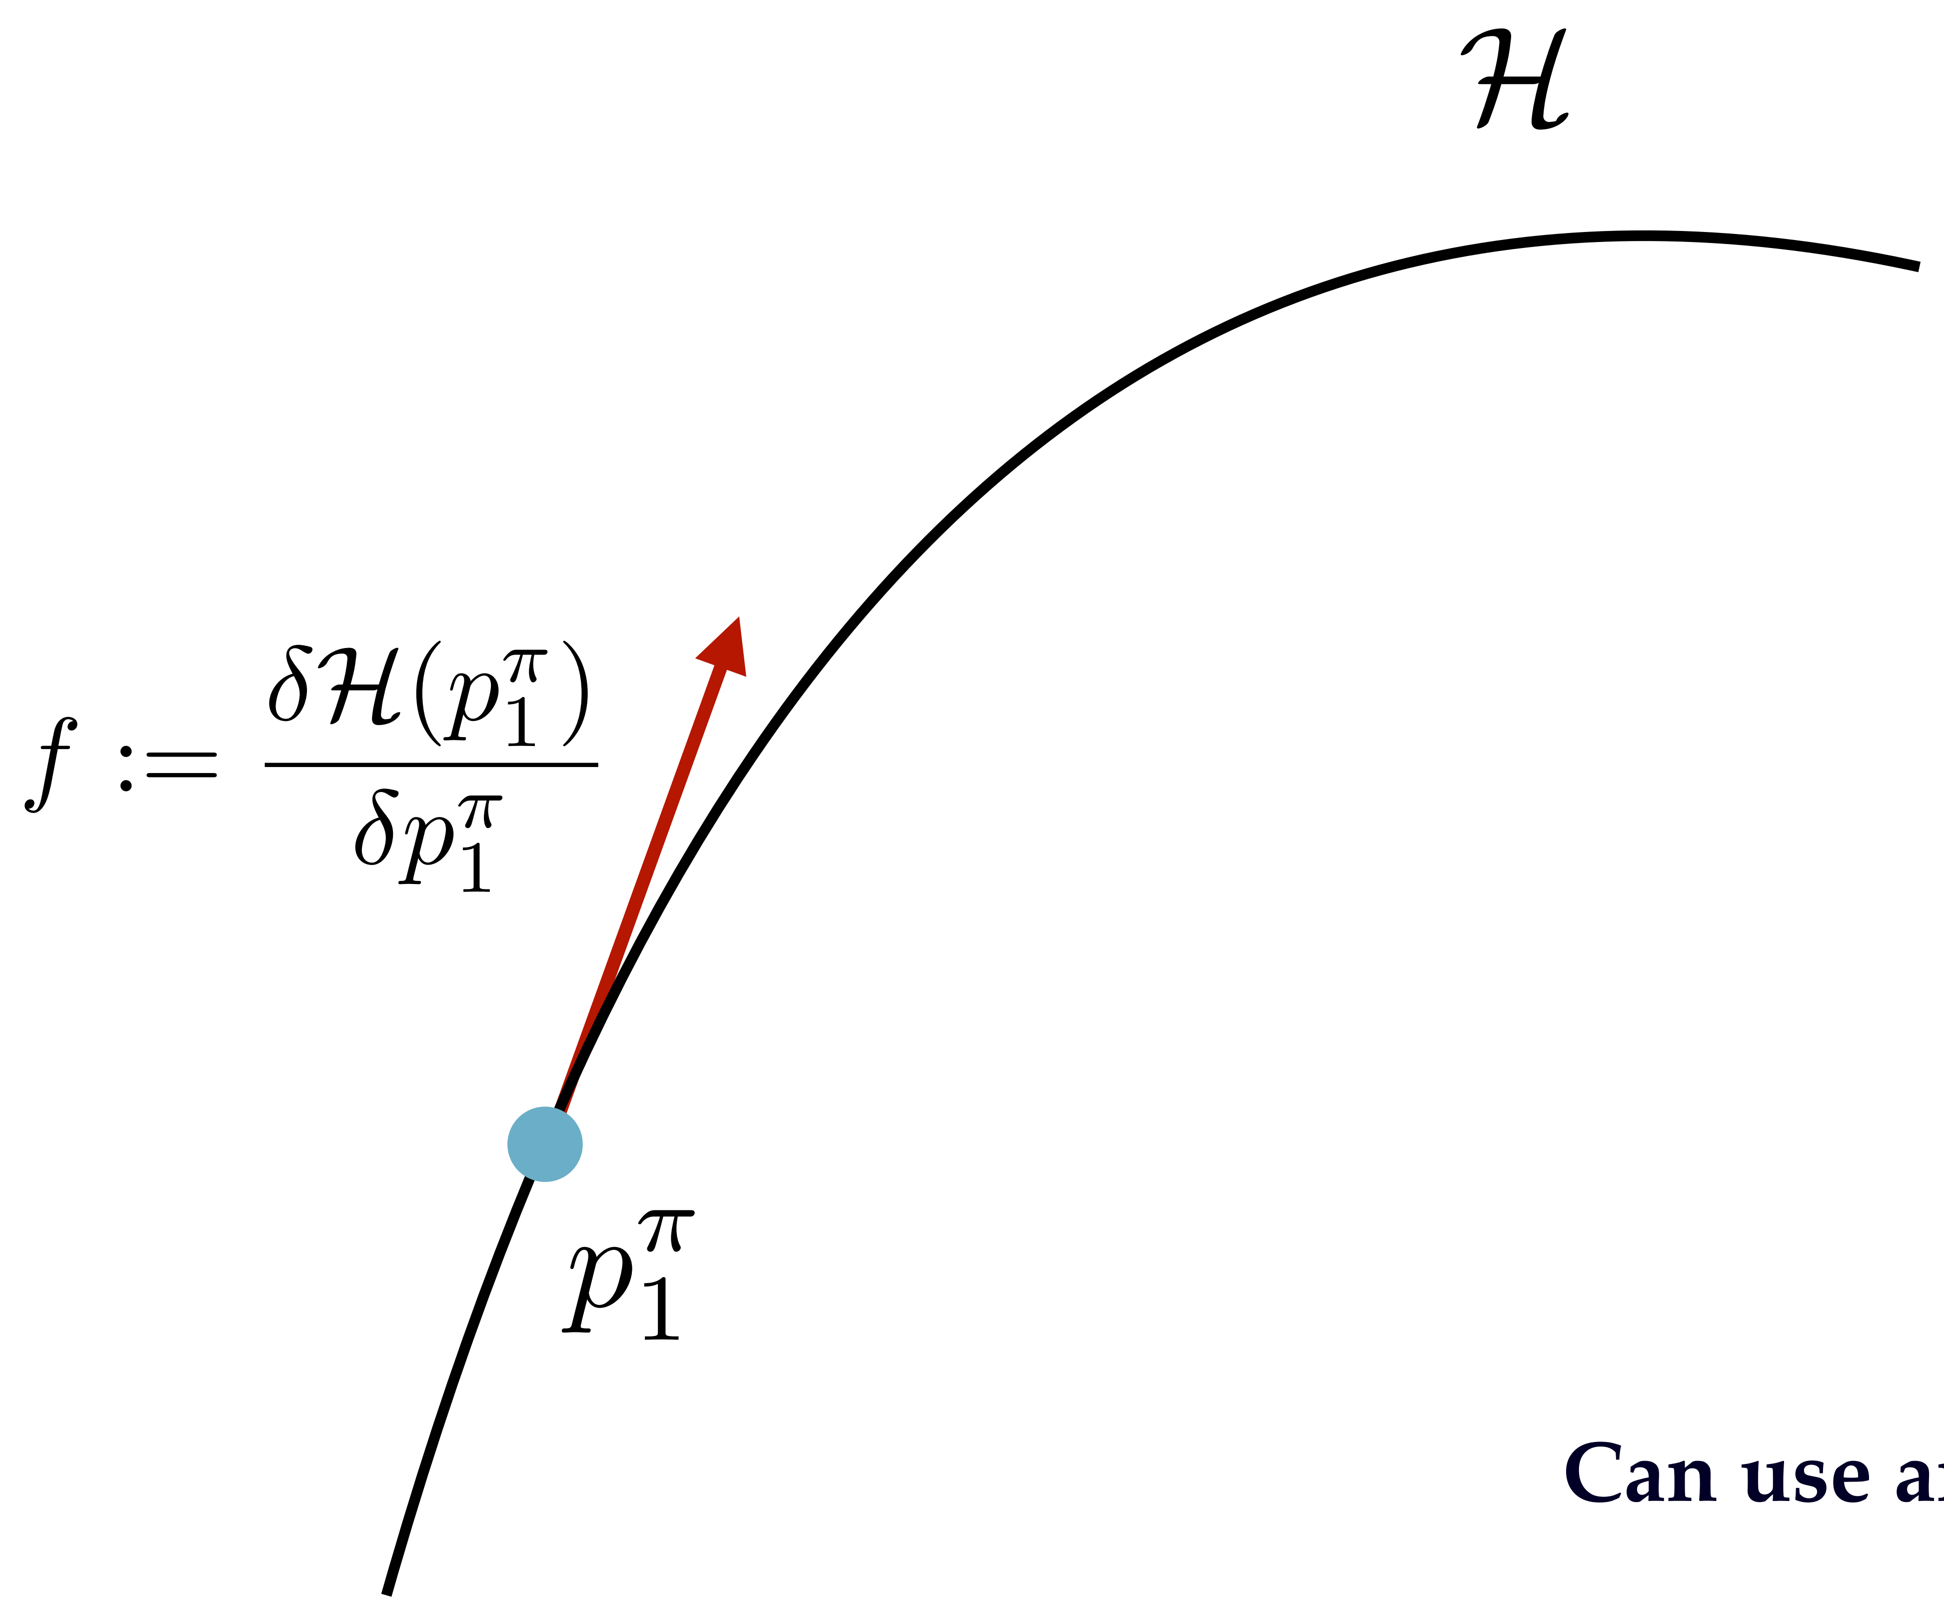

Entropy First Variation (Surprise)

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

surprise as **intrinsic reward function**

Closed-form Gradient of Surprise

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

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

Can use any gradient-based reward-guided FT method.

Key: linearizing entropy leads to **intrinsic reward** f **(i.e. surprise)** for standard **RL fine-tuning**

Flow Density Control (FDC)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \ldots, K$:

Set $\nabla f_k := \nabla \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}} = -(1+\alpha) s_1^{\pi_{k-1}} - s_1^{\pi_{\text{pre}}}$

Fine-tune π_k via standard reward-guided fine-tuning:

$$\pi_k \longleftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi} [f_k(x)] - \eta_k \mathbf{KL}(p_1^\pi \,\|\, p_1^{k-1})$$

Return $\pi := \pi_K$

Takeaway

1. **Data Debiasing | Mode Discovery | Exploration** can be tackled via **KL-regularized entropic expansion**
2. **KL-regularized entropic expansion** can be reduced to **RL via negative score as reward gradient**

Data De-Biasing or Mode Discovery via FDC

Pre-trained generative model

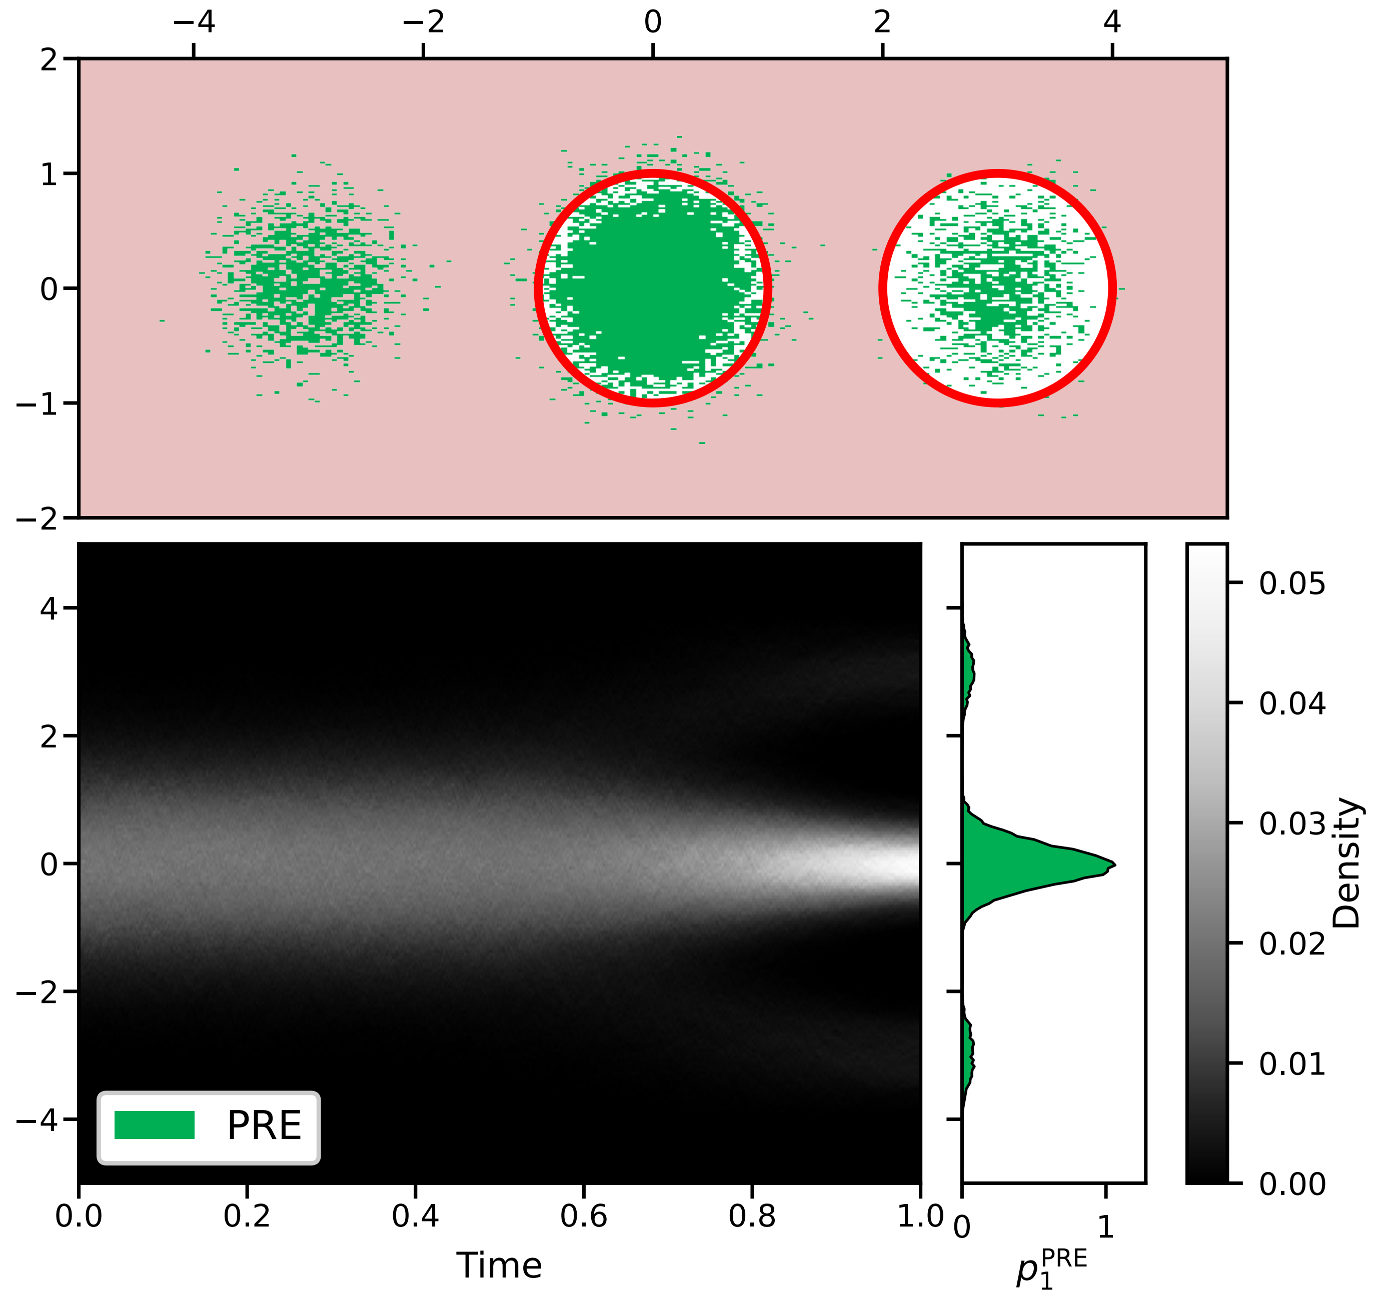

$$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 1$$

Data De-Biasing or Mode Discovery via FDC

Pre-trained generative model

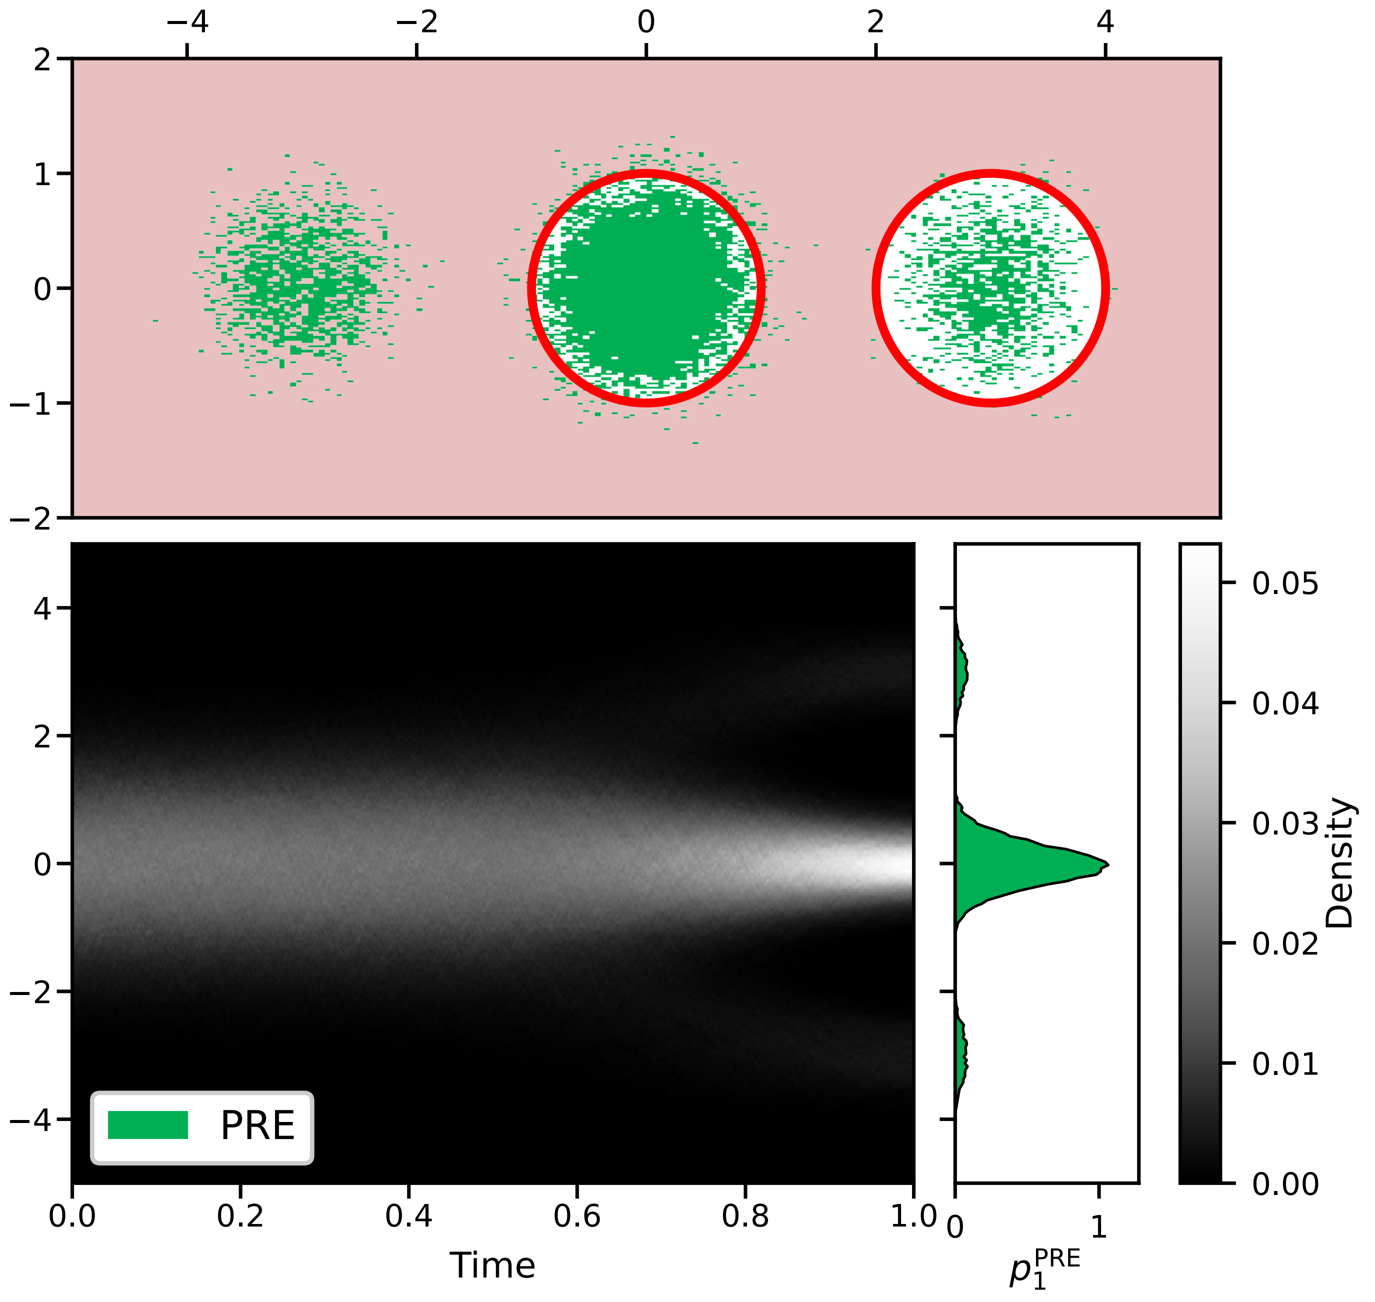

$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 1$

FDC

Expanded generative model

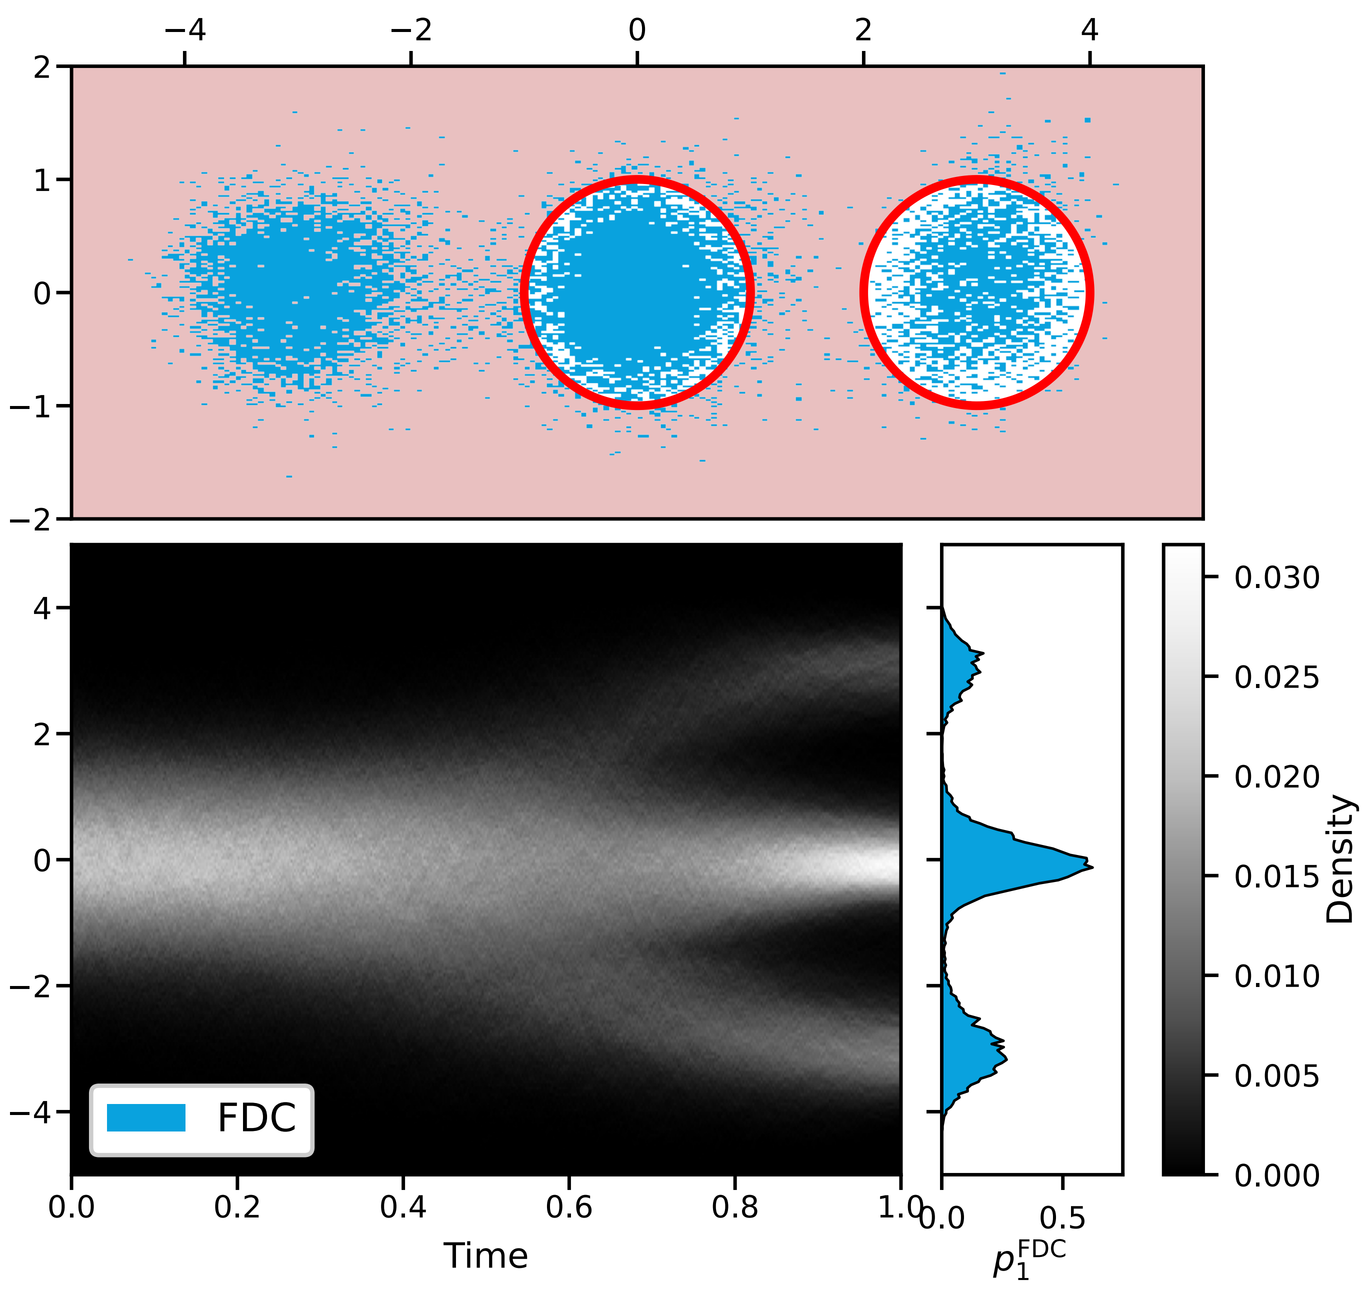

$\mathbb{P}(\text{sample } x \text{ in main mode}) \approx 0.5$

Mode Discovery via FDC: Molecule and Protein Design

GEOM-Drugs

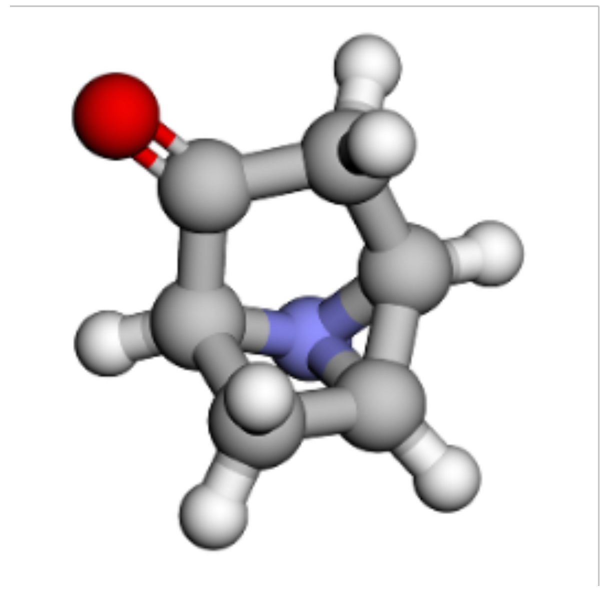

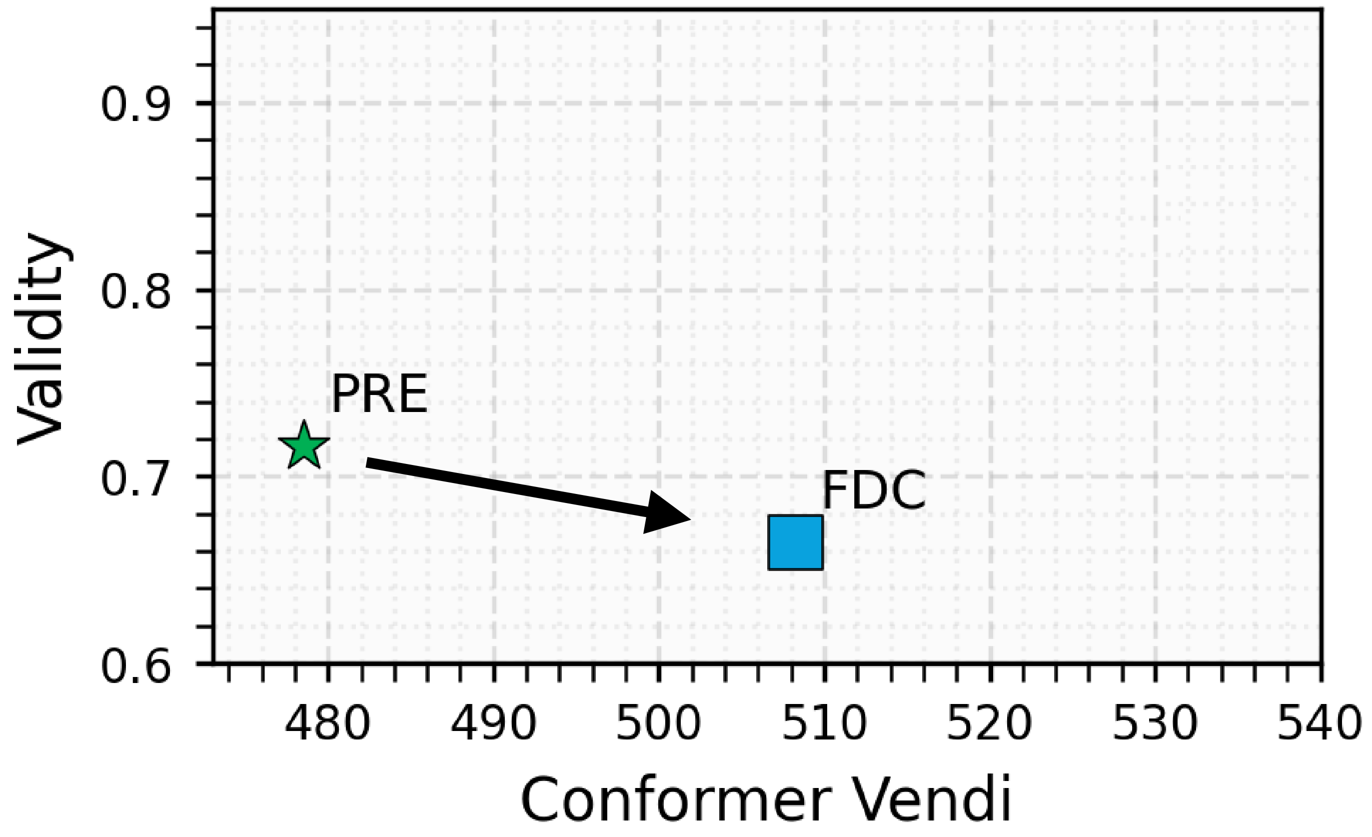

FDC achieves **higher conformer diversity** of drug-like molecules

FDC: Theoretical Guarantees

Assumption (Informal, Exact estimation and optimization).

1. Exact score estimation, i.e., $s^{pre}(\cdot, 1) = \nabla_x \log p_1^{pre}$

2. Reward fine-tuning is solved exactly.

FDC: Theoretical Guarantees

Assumption (Informal, Exact estimation and optimization).

1. Exact score estimation, i.e., $s^{pre}(\cdot, 1) = \nabla_x \log p_1^{pre}$

2. Reward fine-tuning is solved exactly.

↓

Theorem (Informal, One-step convergence). Fine-tuning a pre-trained model π^{pre} for entropy-maximization via **FDC** leads to a model π such that:

$$\mathcal{H}(p_1^*) - \mathcal{H}(p_1^\pi) = 0$$

where $p_1^* := p_1^{\pi^*}$ is the marginal distribution induced by the optimal exploratory policy $\pi^* \in \arg\max_{\pi \in \Lambda} \mathcal{H}(p_1^\pi)$ with $\Lambda = \{\pi : p_1^\pi \in \mathbb{P}(\Omega_{pre})\}$.

This talk:

Foundations of Generative Discovery Beyond the Data

Distributional Generative Optimization

Part I

Tail-Aware Reward Adaptation

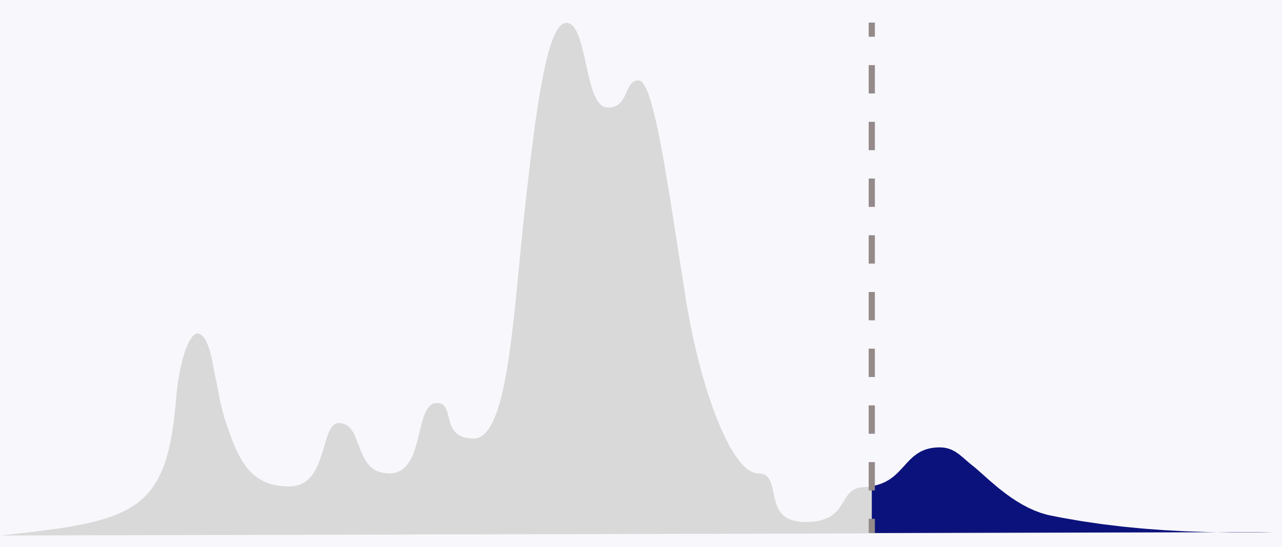

$\mathrm{CVaR}^{f}_{1-\beta}(p_1^{\pi})$

Part II

Data Debiasing and Hidden Mode Discovery

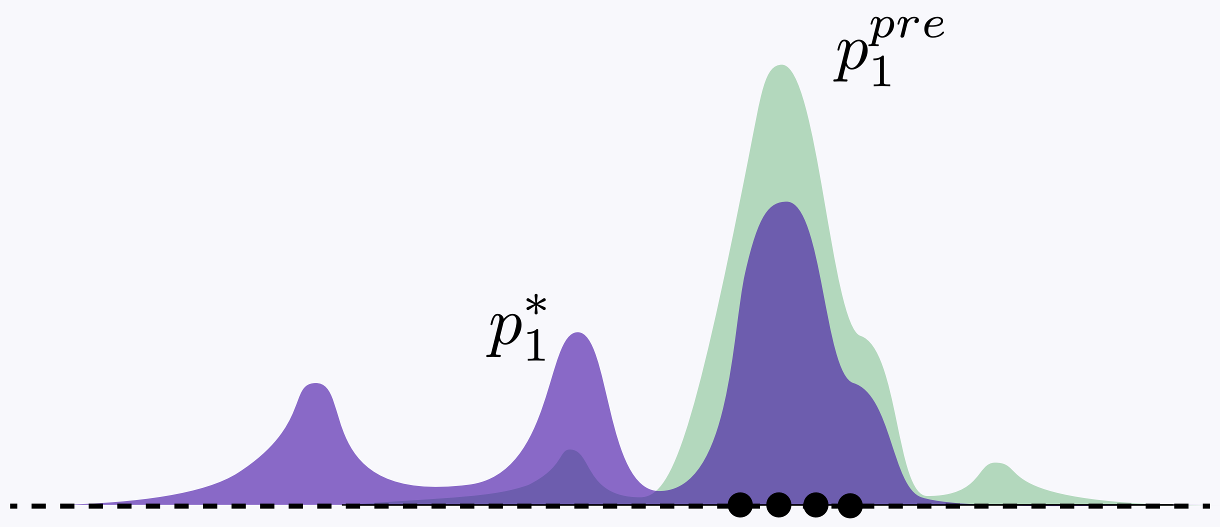

Part III

Out-of-Distribution Flow Modeling

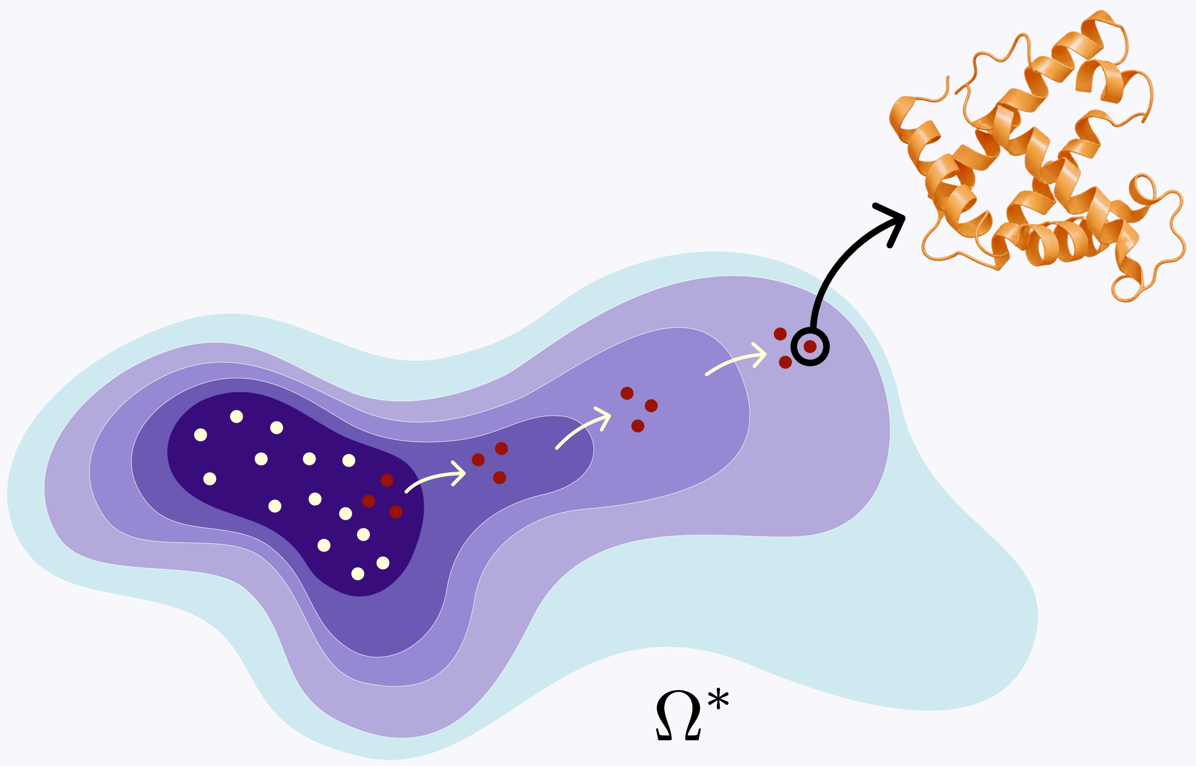

Flow Density Control for Distributional Fine-Tuning

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

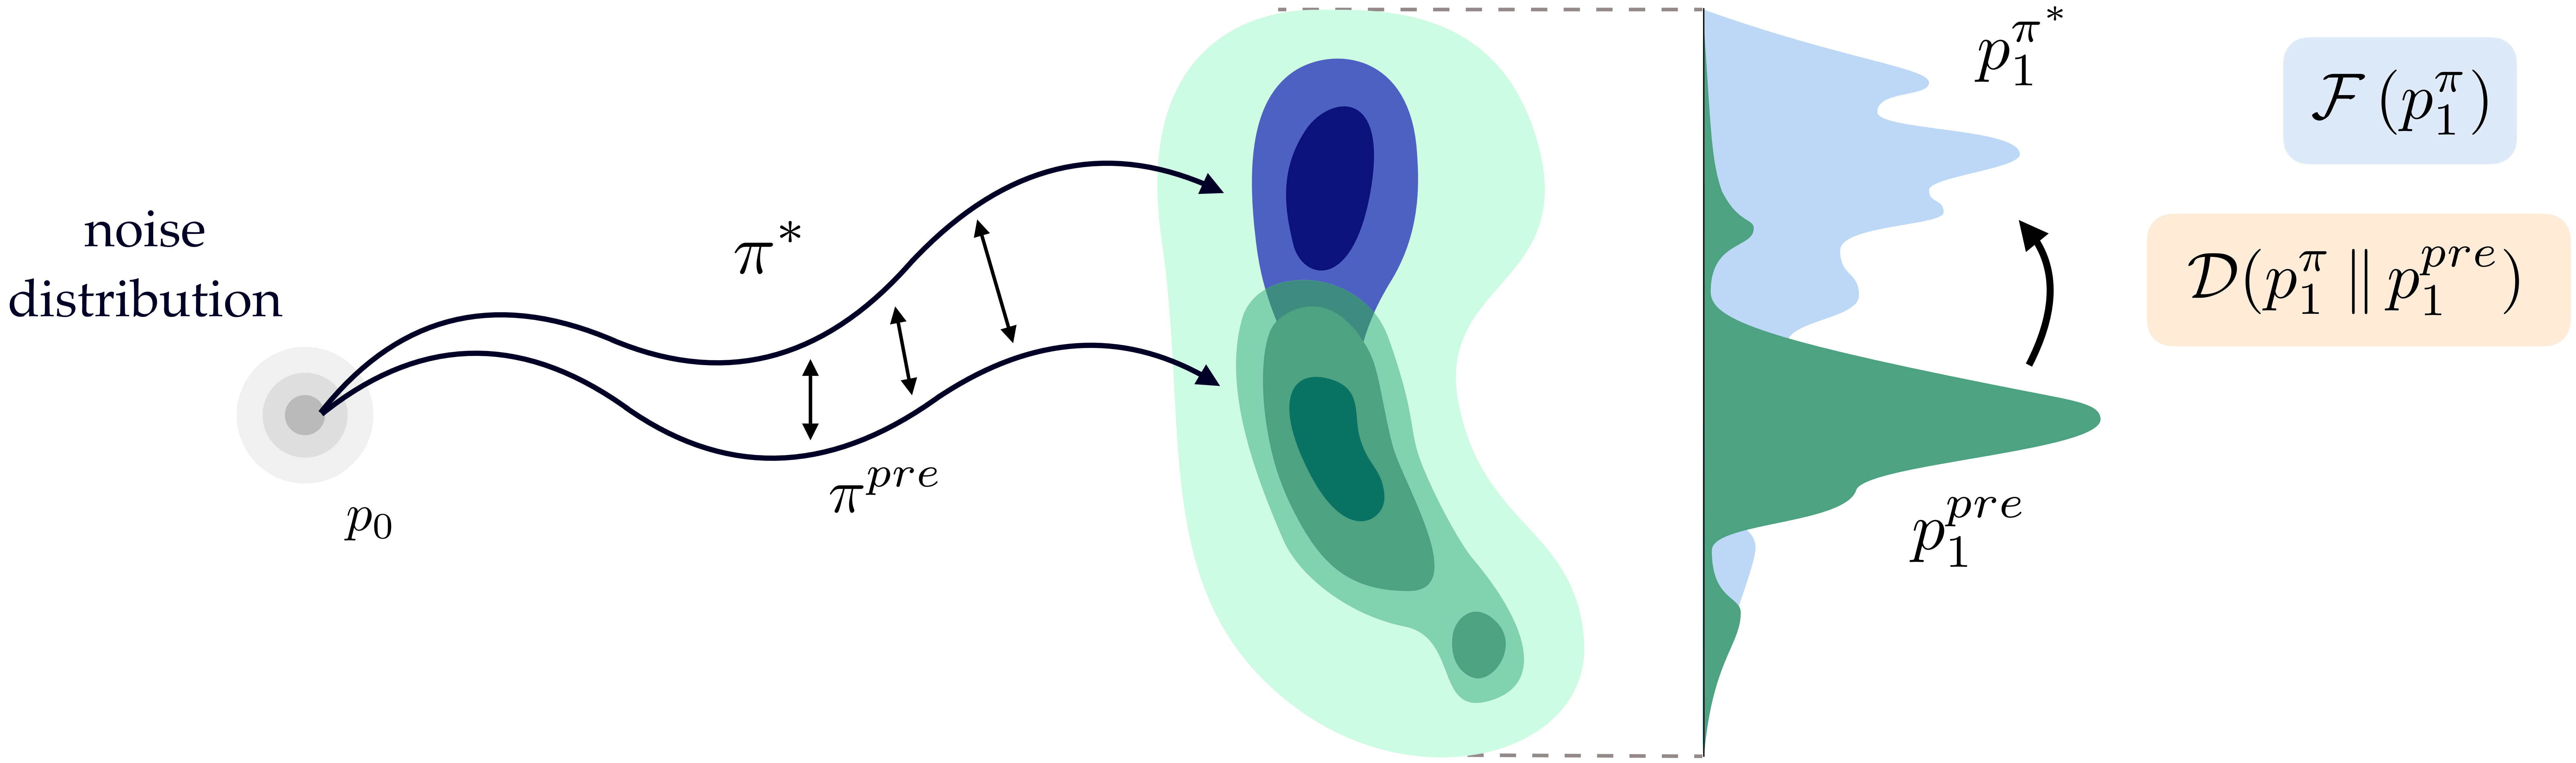

How to represent new task? $\mathcal{F}$

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

Distributional Generative Optimization

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

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	Non-Linear GO: General
Reward optimization [14, 55]	$\mathbb{E}_{x\sim p^\pi}[r(x)]$	✓	✓	✓
Manifold Exploration [12] Gen. model de-biasing	$\mathcal{H}(p^\pi) := -\mathbb{E}_{x\sim p^\pi}[\log p^\pi(x)]$	✗	✓	✓
Risk-averse optimization	$\text{CVaR}^r_\beta(p^\pi) := \mathbb{E}_{x\sim p^\pi}[r(x) \mid r(x) \leq \text{q}^r_\beta(p^\pi)]$	✗	✓	✓
	$\mathbb{E}_{x\sim p^\pi}[r(x)] - \mathbb{V}\text{ar}(p^\pi)$	✗	✗	✓
Novelty-seeking optimization	$\text{SQ}^r_\beta(p^\pi) := \mathbb{E}_{x\sim p^\pi}[r(x) \mid r(x) \geq \text{q}^r_\beta(p^\pi)]$	✗	✗	✓
Optimal Experiment Design	$\text{s}\left(\mathbb{E}_{x\sim p^\pi}[\Phi(x)\Phi(x)^\top - \lambda\mathbb{I}]\right)$ $\text{s}(\cdot) \in \{\log\det(\cdot), -\text{Tr}(\cdot)^{-1}, -\lambda_{max}(\cdot)\}$	✗	✓	✓
Diverse modes discovery	$-\mathbb{E}_z[D_{KL}(p^{\pi,z} \| \mathbb{E}_k p^{\pi,k})]$	✗	✗	✓
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]	$D_{KL}(p^\pi \| p^{pre}) = \int p^\pi(x)\log\frac{p^\pi(x)}{p^{pre}(x)}dx$	✓	✓	✓
Rényi divergences	$D_\beta(p^\pi \| p^{pre}) := \frac{1}{\beta-1}\log\int(p^\pi(x))^\beta(p^{pre})^{1-\beta}dx$	✗	✗	✓
Optimal Transport distances	$W_p(p^\pi \| p^{pre}) := \inf_{\gamma\in\Gamma(p^\pi,p^{pre})}\mathbb{E}_{(x,y)\sim\gamma}[d(x,y)^p]^{\frac{1}{p}}$	✗	✗	✓
Maximum Mean Discrepancy	$\text{MMD}_k(p^\pi \| p^{pre}) := \|\mu_{p^\pi} - \mu_{p^{pre}}\|,\ \mu_p := \mathbb{E}_{x\sim p}[k(x,\cdot)]$	✗	✓	✓

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

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

$$\downarrow$$

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

Flow Density Control: Theoretical Guarantees

Distributional Generative Optimization

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

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

<u>concave, idealized assumptions</u>, 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)

<u>general (e.g., non-concave) assumptions:</u>

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

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

(Asymptotic convergence via mirror flows)

The Dream: Discovery via Generative Optimization

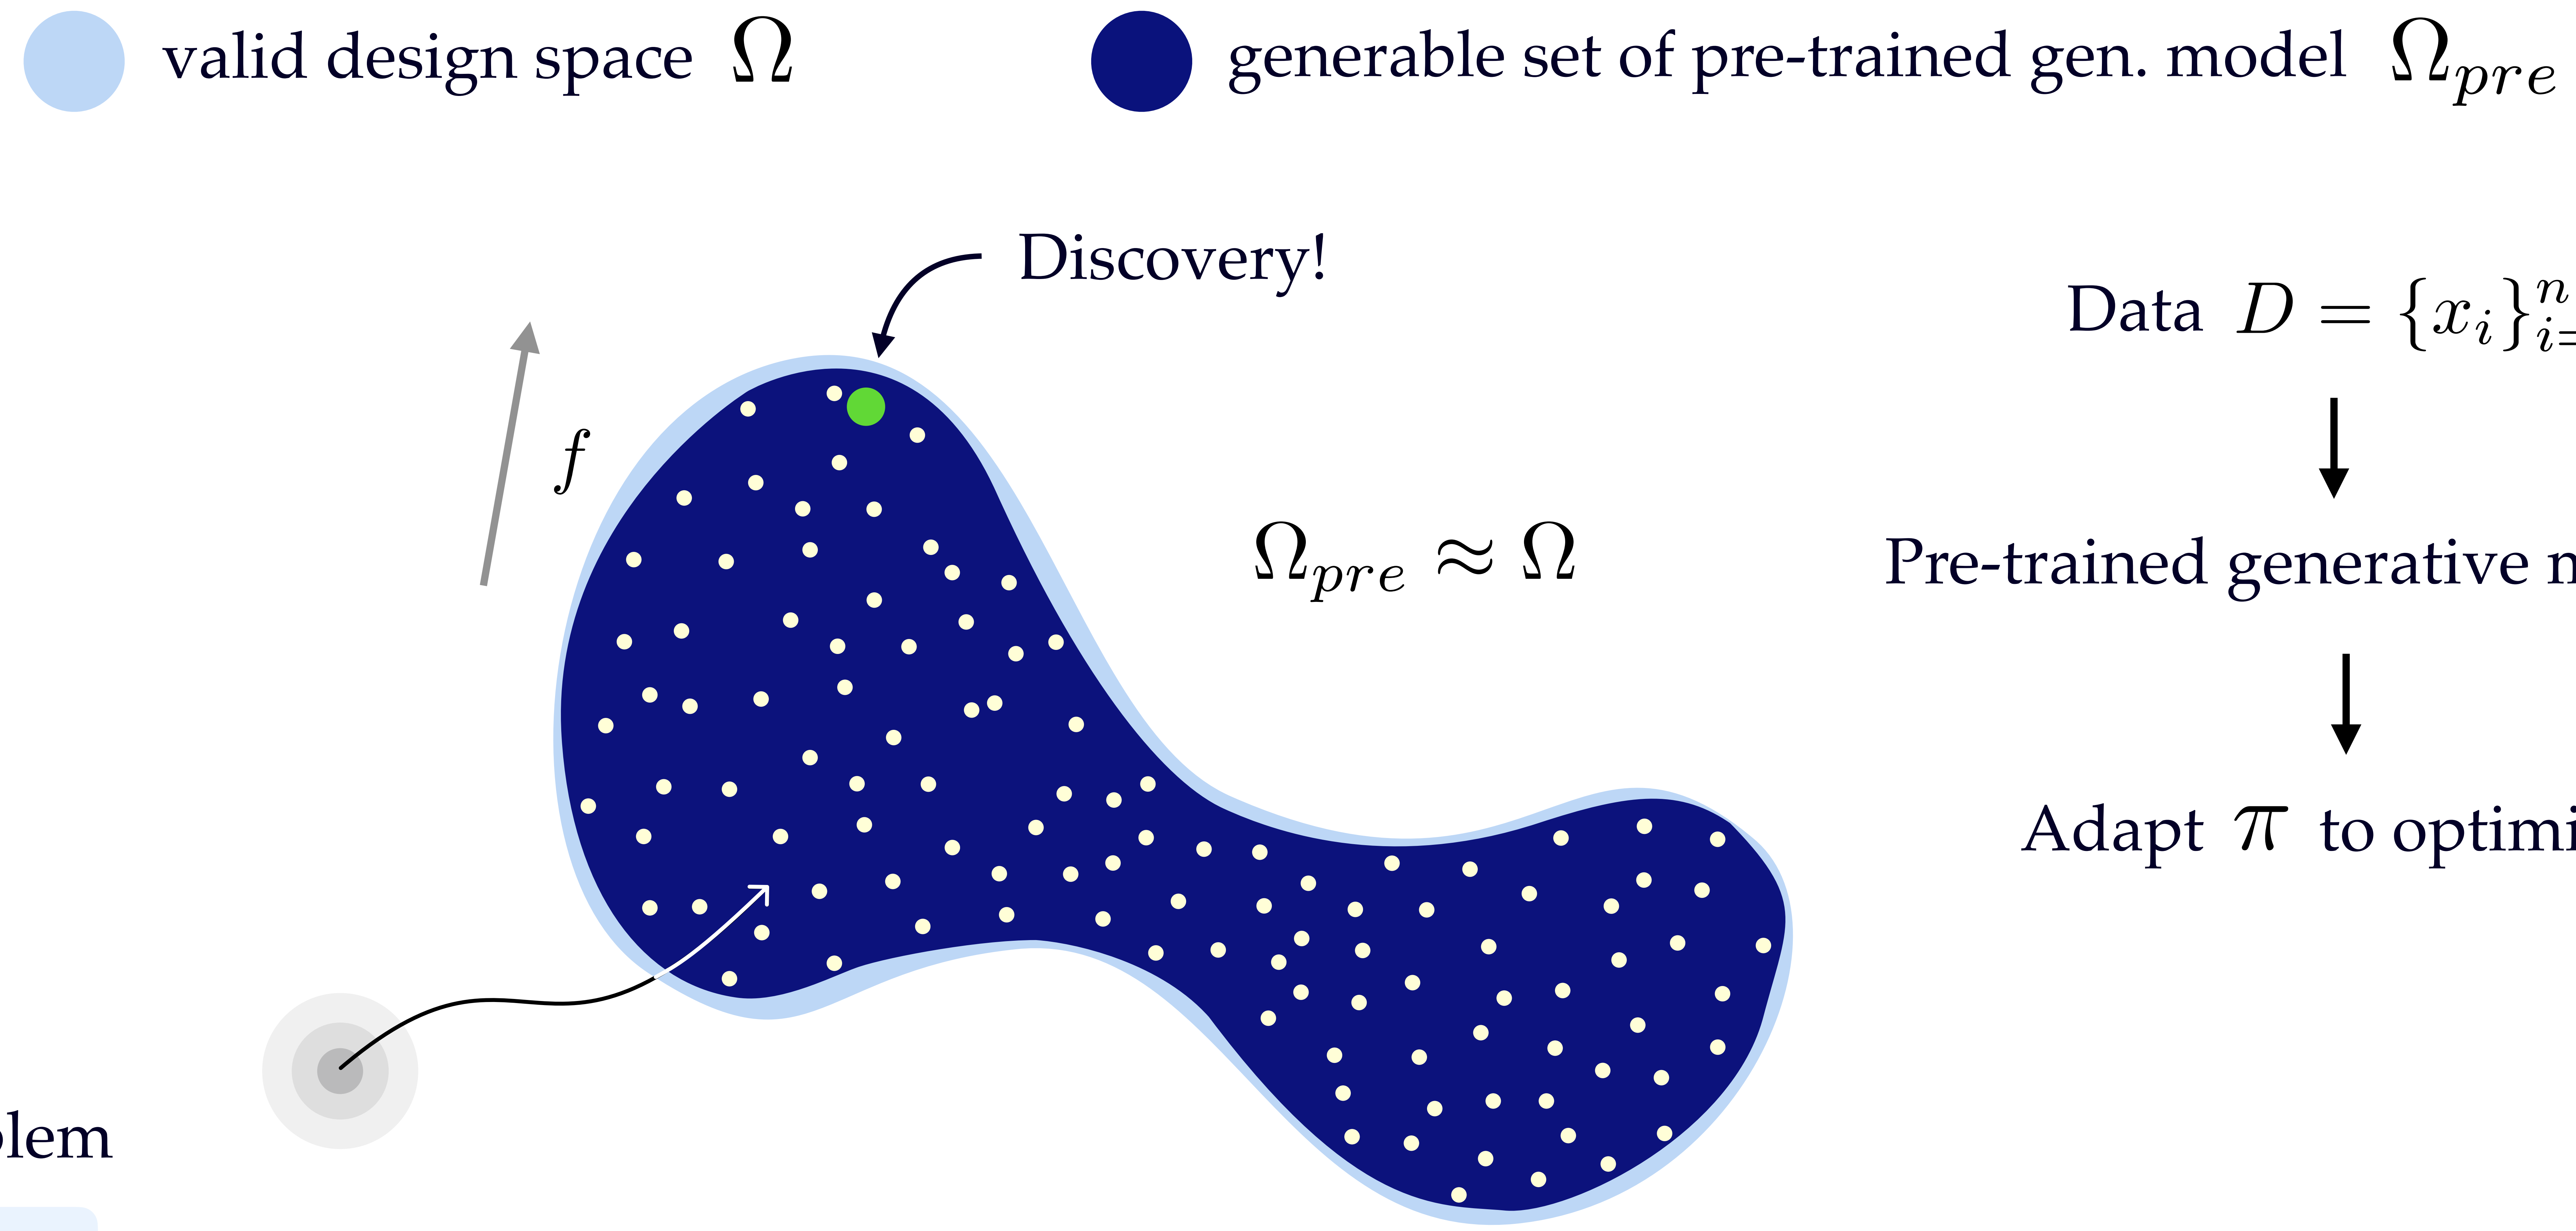

Data $D = \{x_i\}_{i=1}^n$

↓

Pre-trained generative model π

↓

Adapt π to optimize f

Discovery problem

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

The Reality: Generative Optimization Limitation

valid design space Ω

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

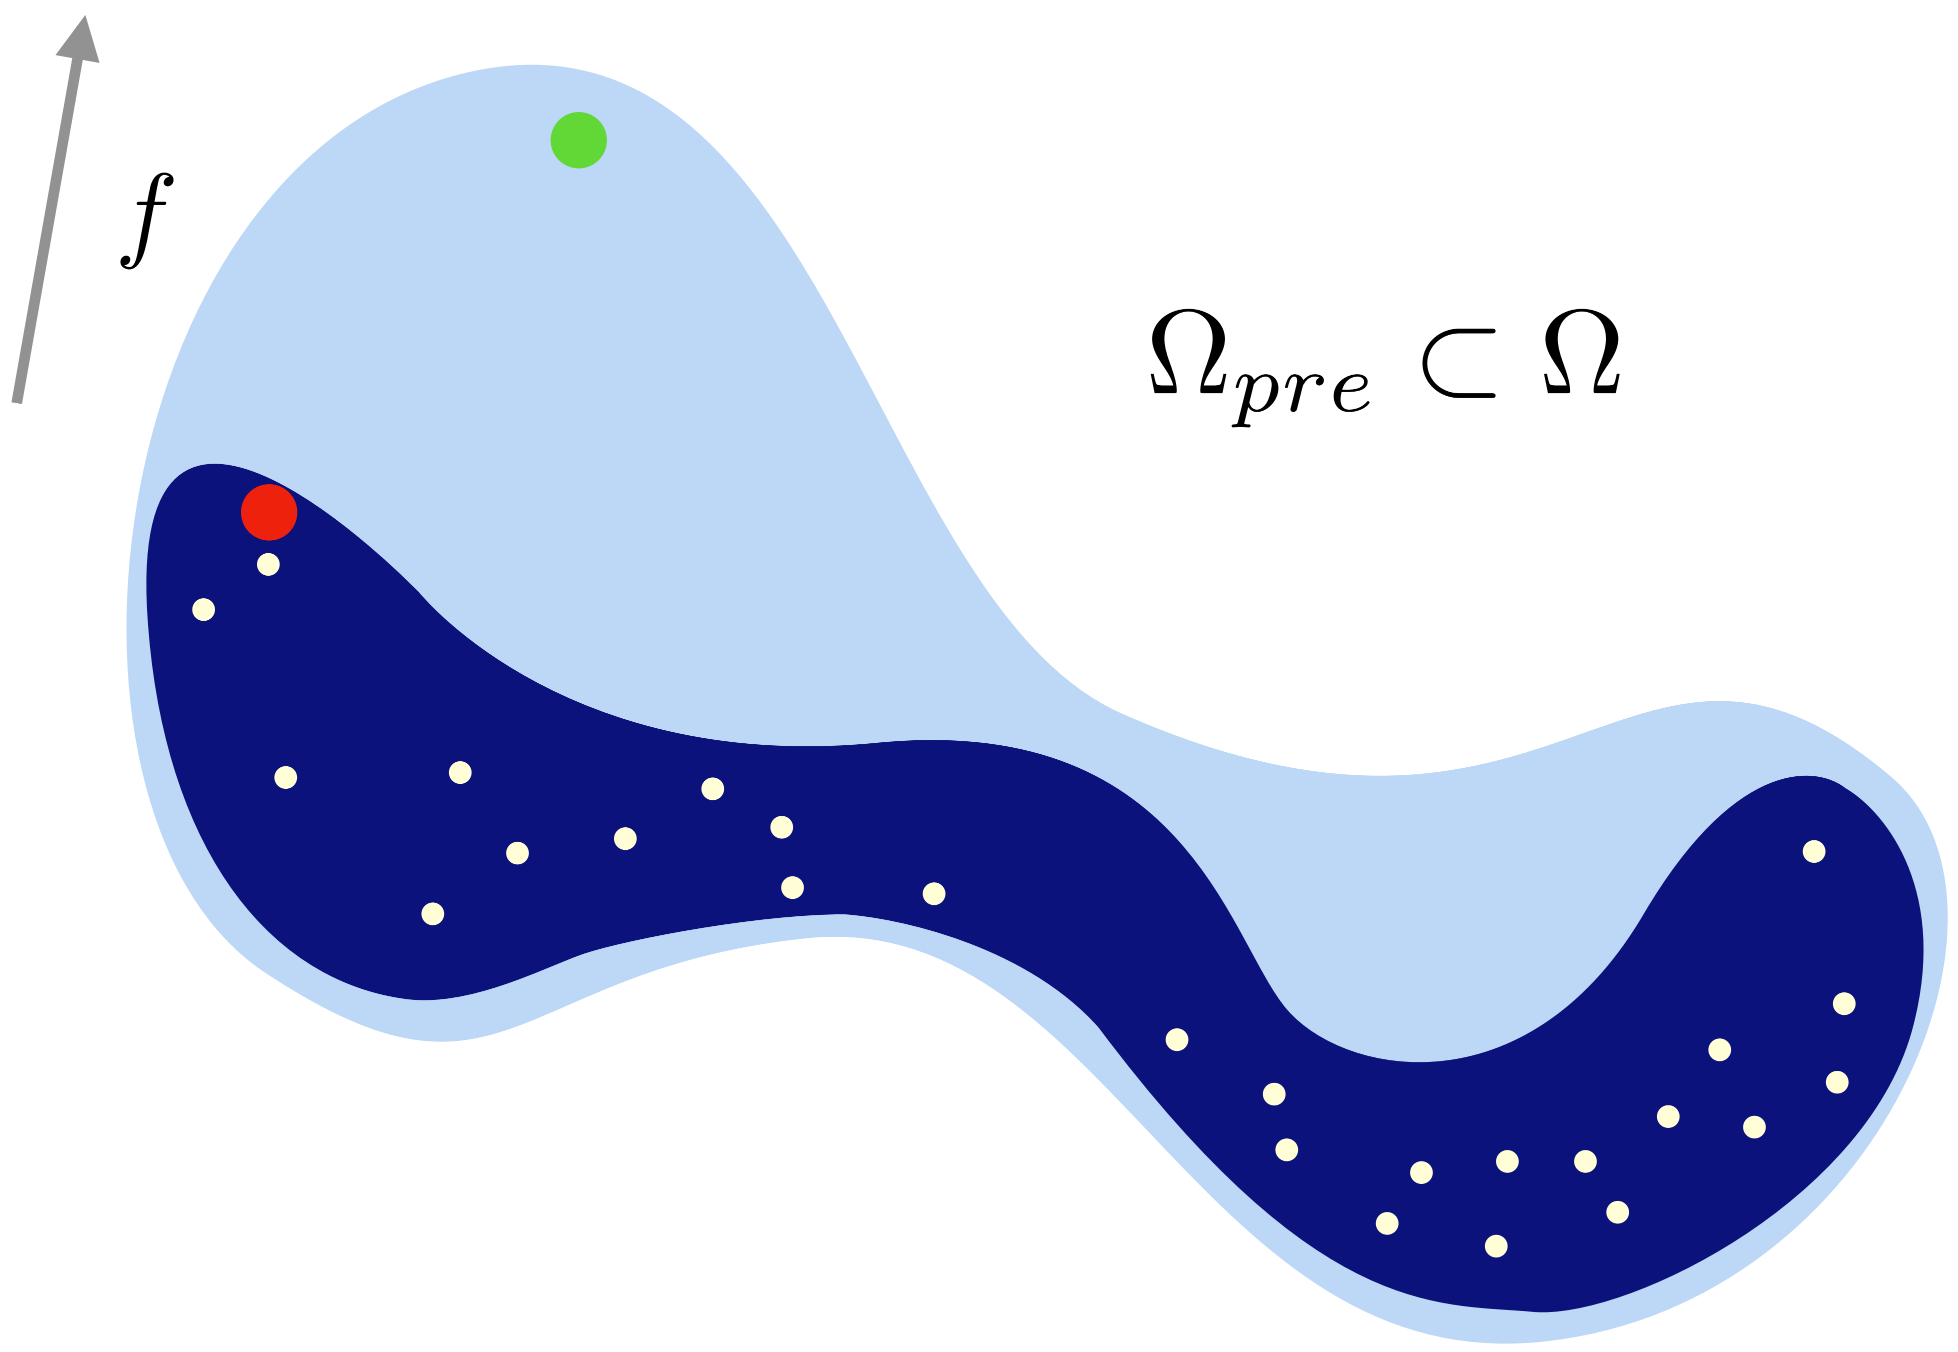

$\Omega_{pre} \subset \Omega$

Generative models generalize locally over training data.

Discovery problem

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

The Reality: Generative Optimization Limitation

valid design space Ω generable set of pre-trained gen. model Ω_{pre}

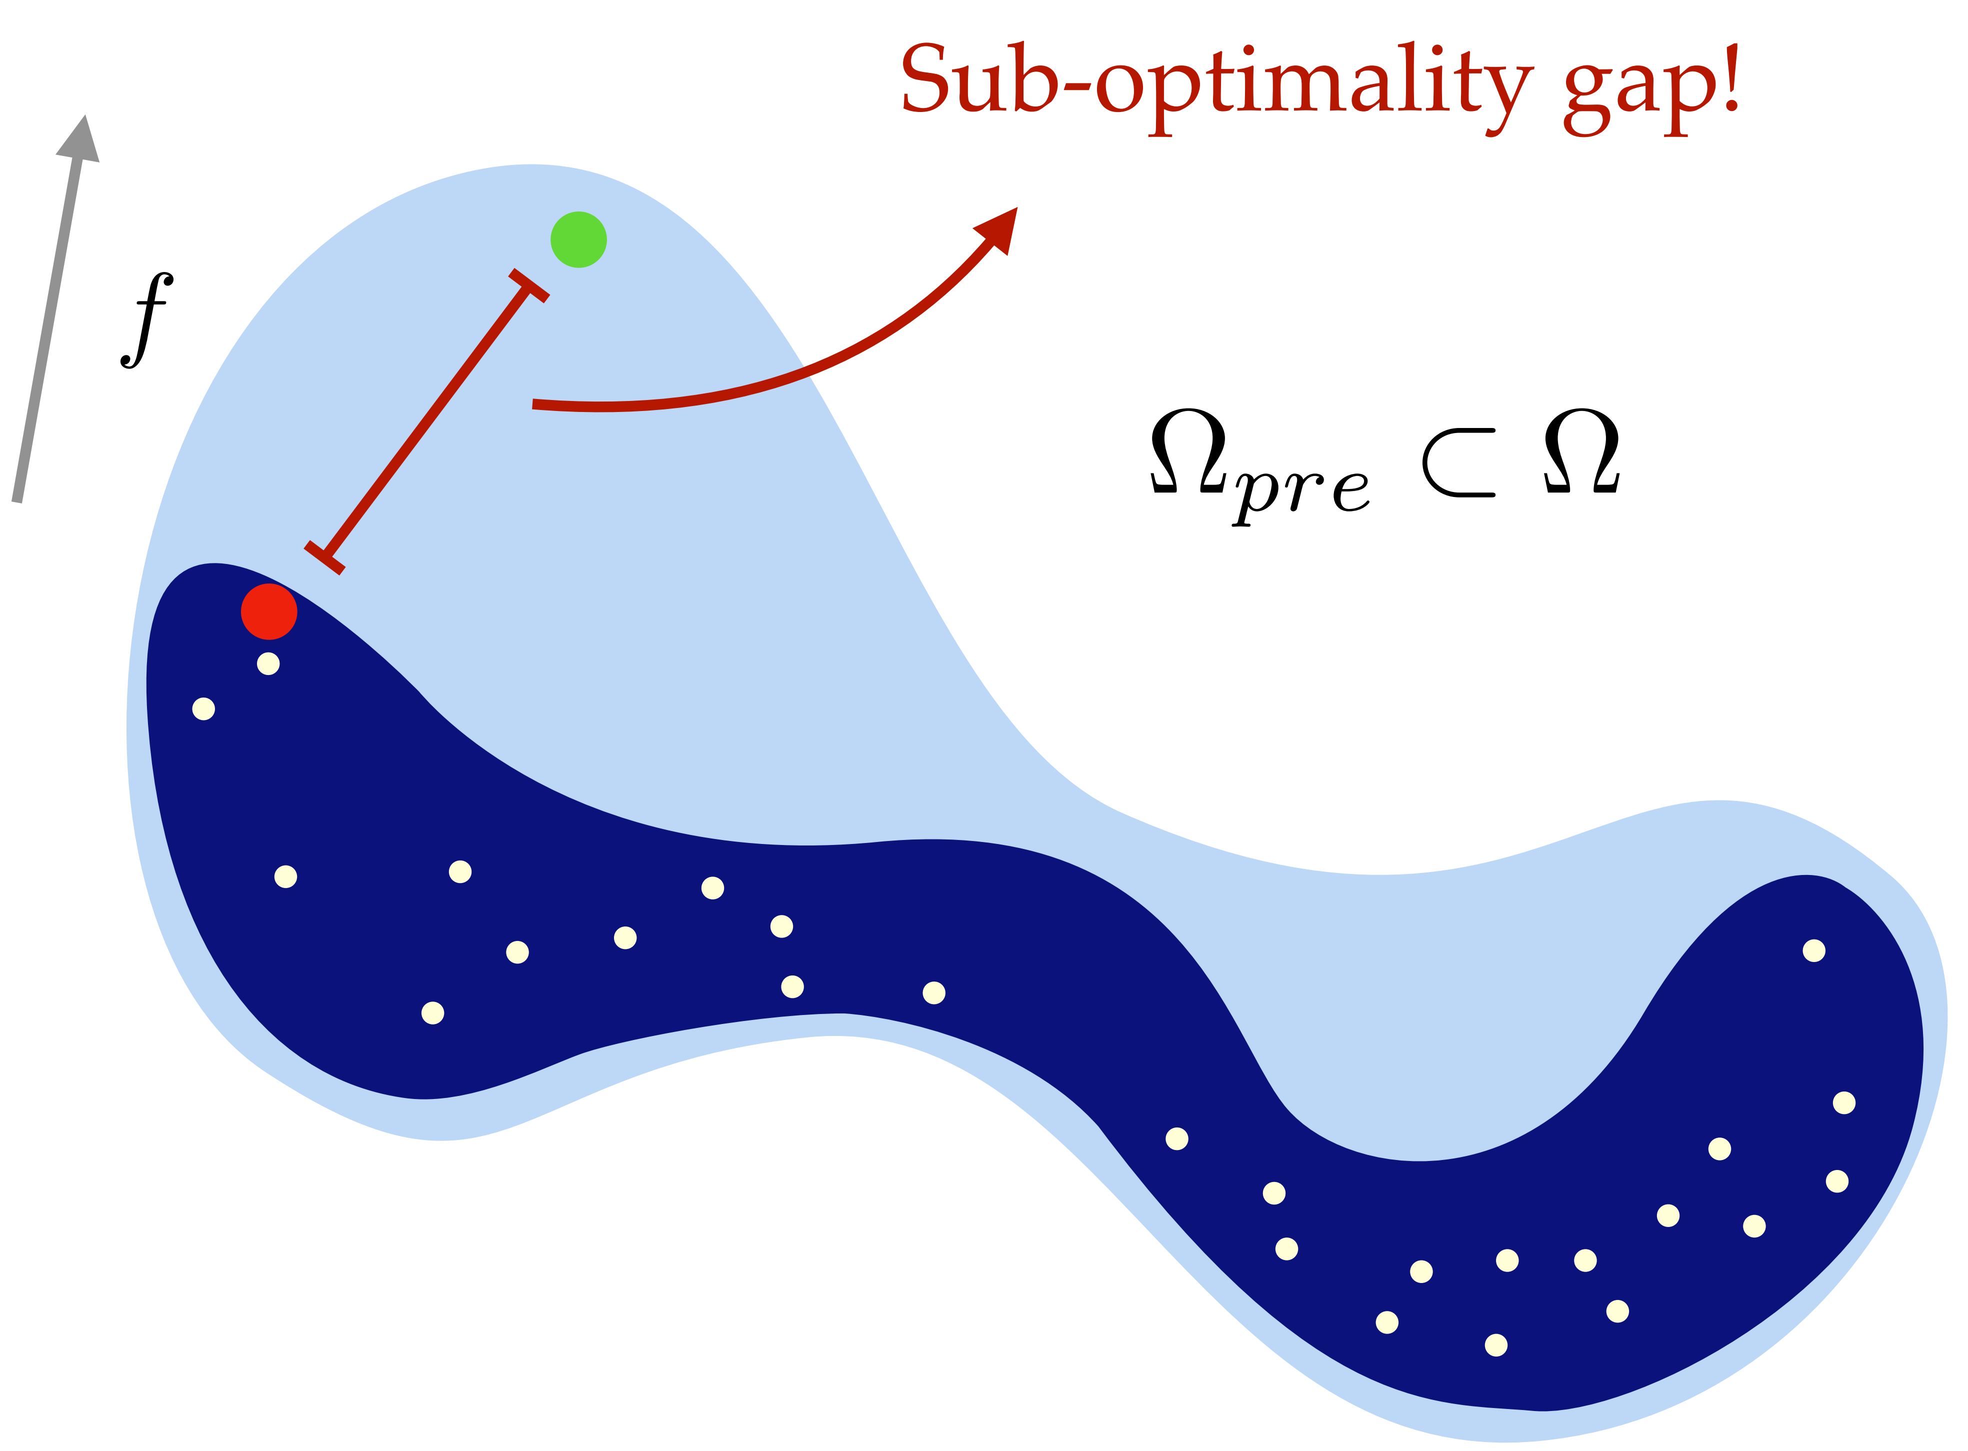

Generative models generalize locally over training data.

Discovery problem

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

Generative discovery entails a new exploration problem

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

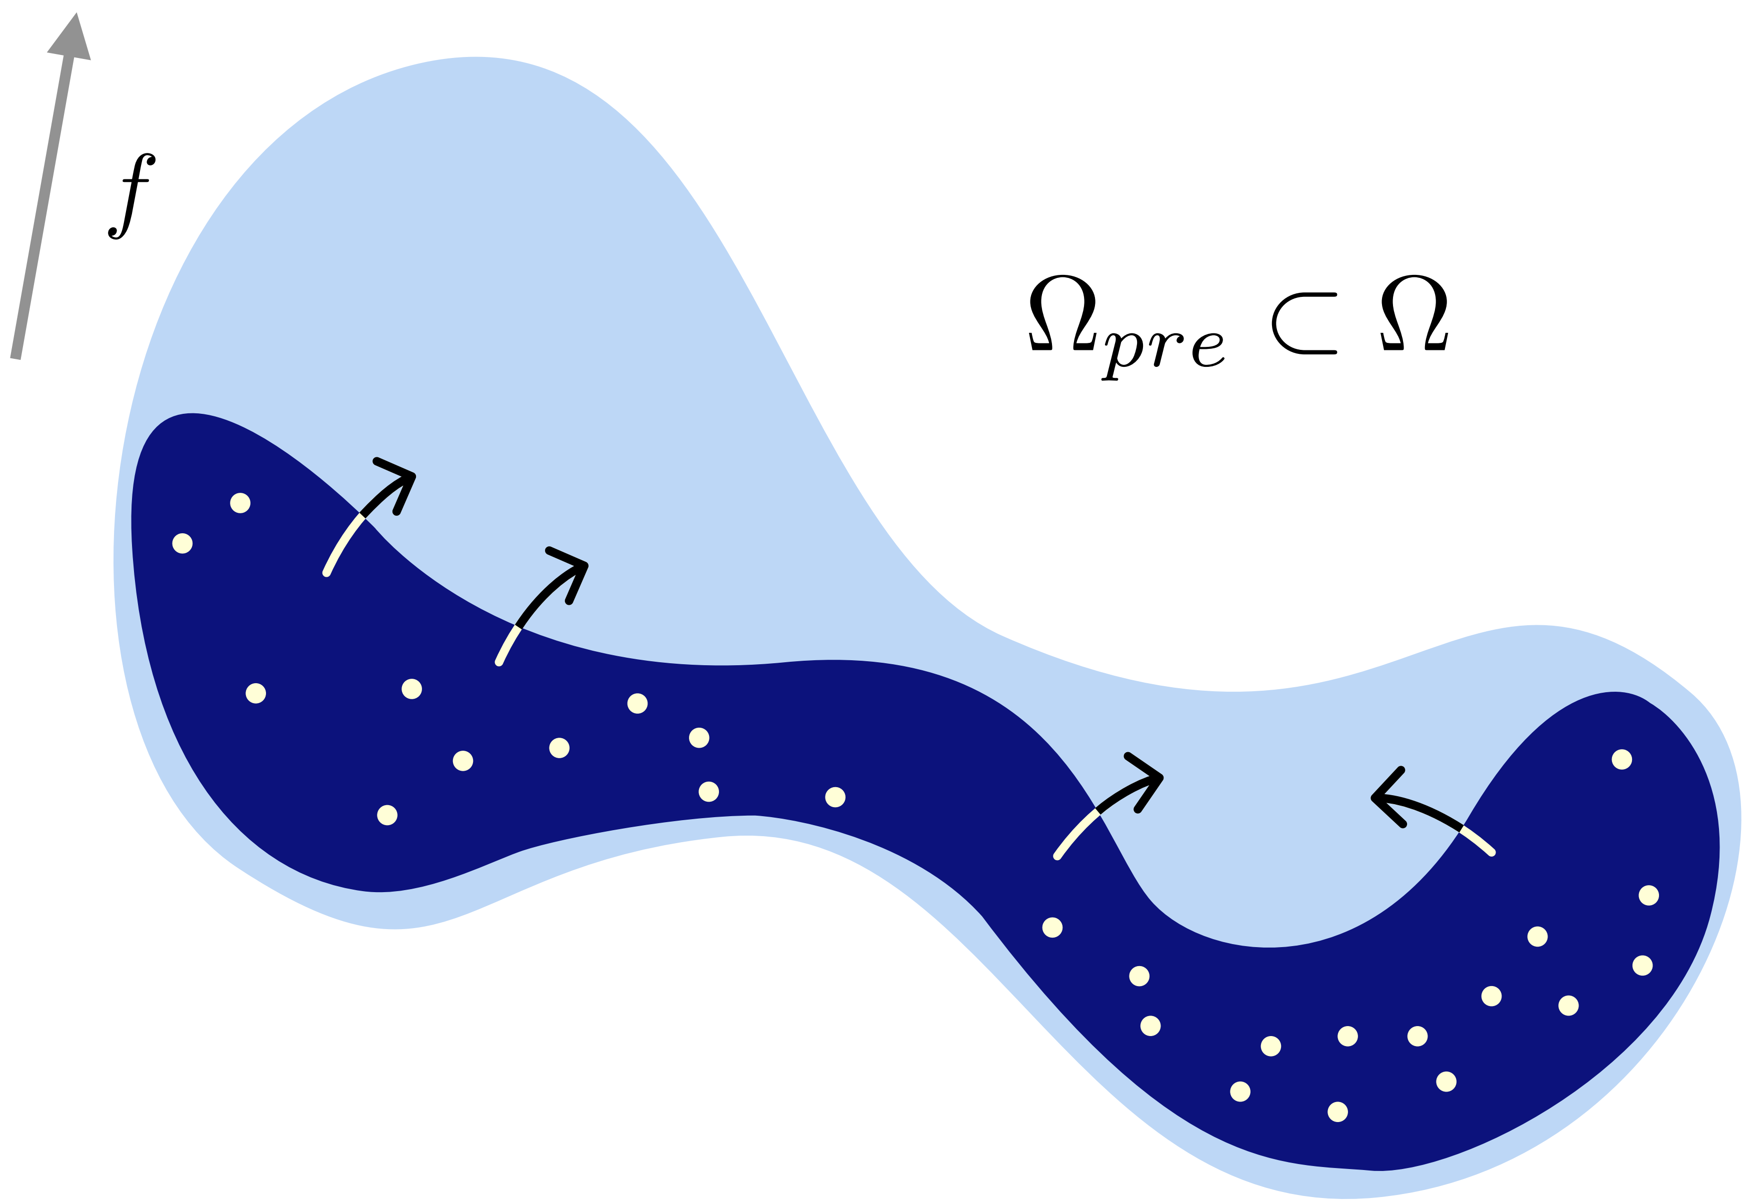

Generative discovery entails a new exploration problem

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

Generative discovery entails a new exploration problem

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

classic optimization
(e.g., convex optimization)

f known

Ω known

Generative discovery entails a new exploration problem

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

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

Generative discovery entails a new exploration problem

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

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

Generative discovery entails a new exploration problem

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

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

Discovery requires the machine to **expand its action space** (i.e., **learn what actions are possible**)

Next: Out-of-Distribution Flow Modeling

valid design space Ω generable set of pre-trained gen. model Ω_{pre}

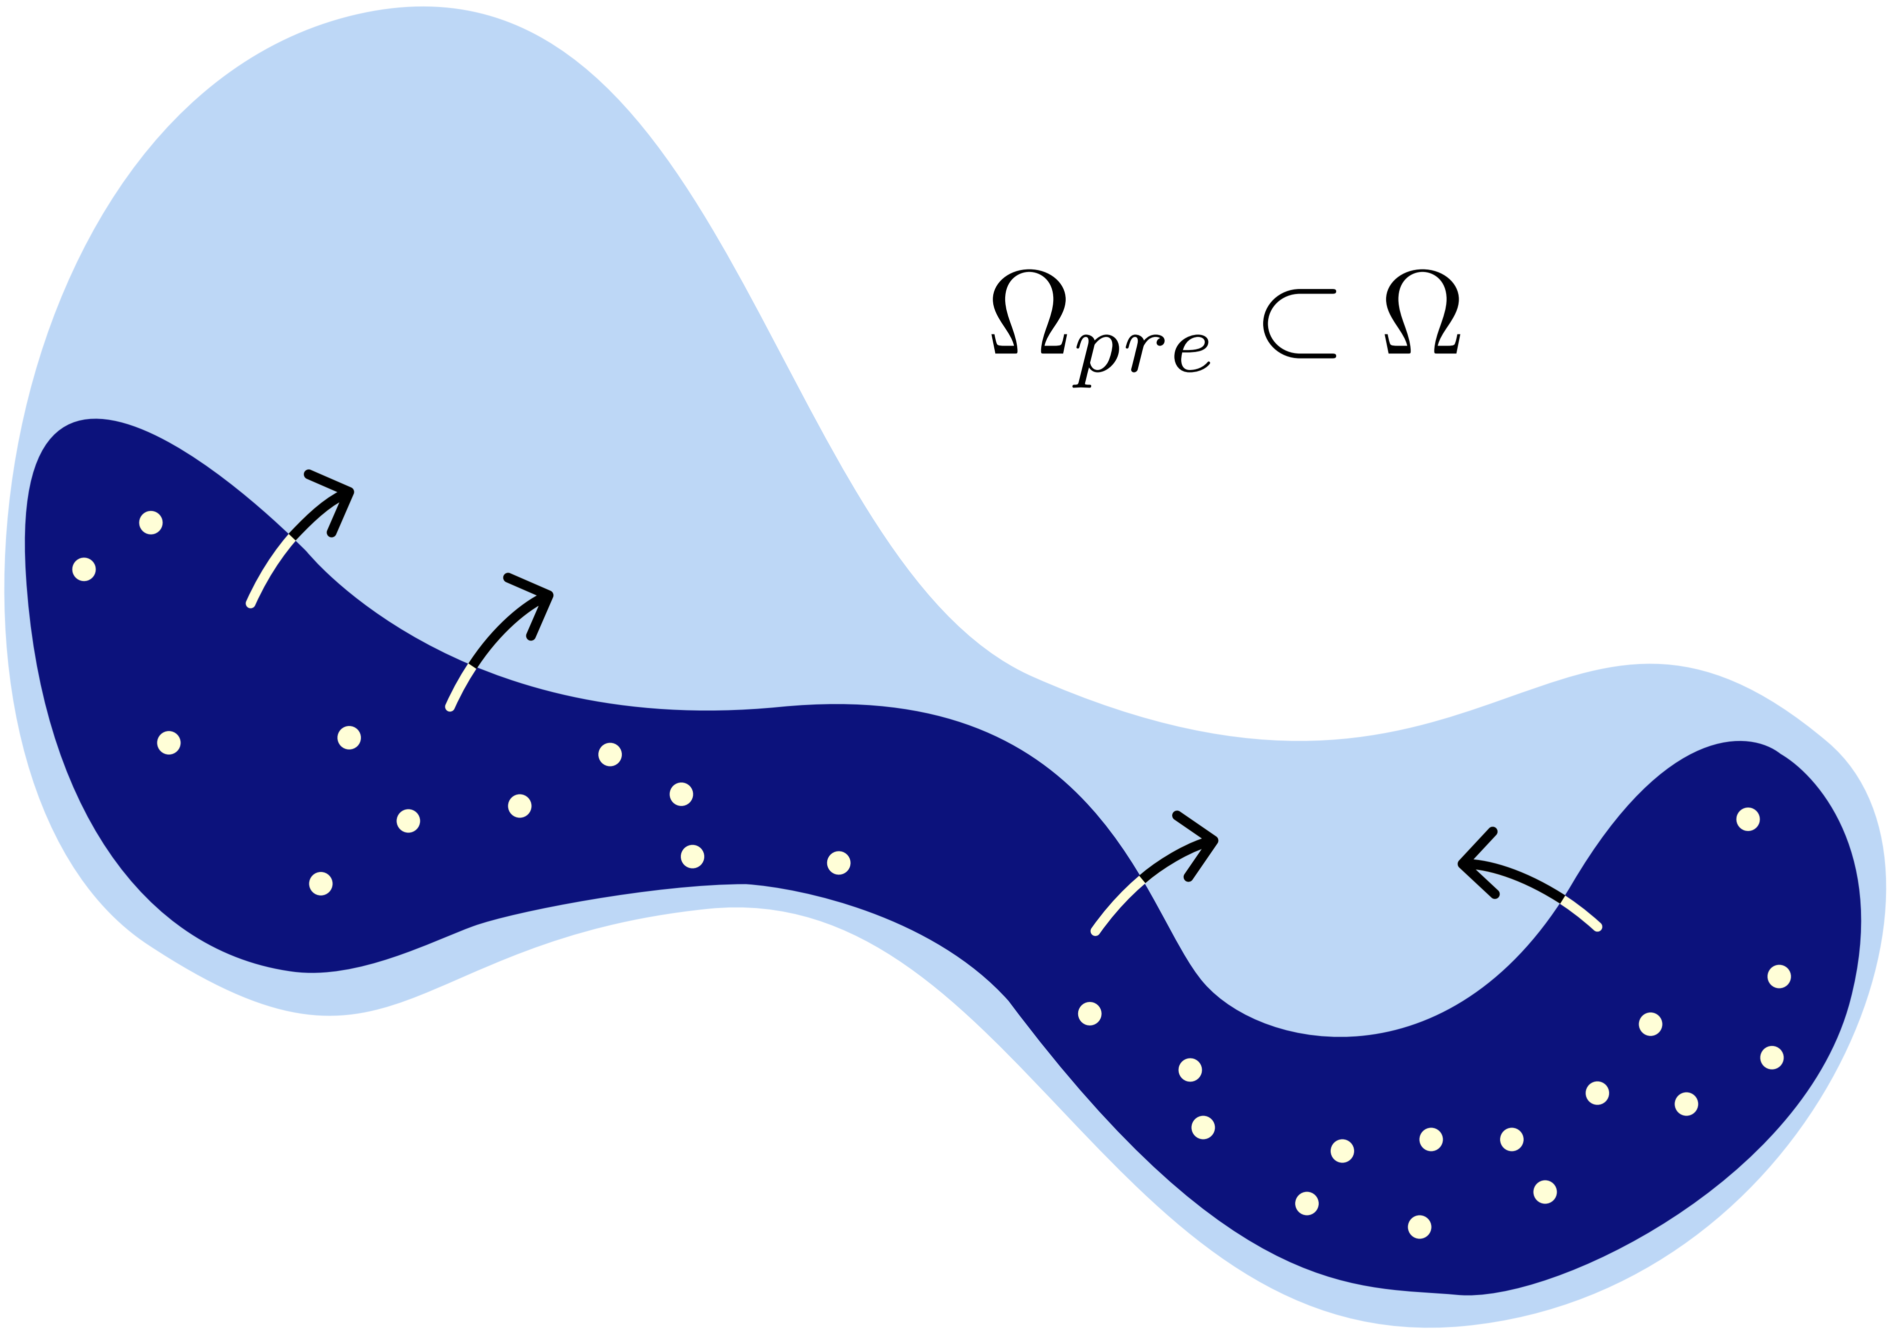

Discovery requires the machine to **expand its action space** (i.e., **learn what actions are possible**)

This talk:

Foundations of Generative Discovery Beyond the Data

Part I

Tail-Aware Reward Adaptation

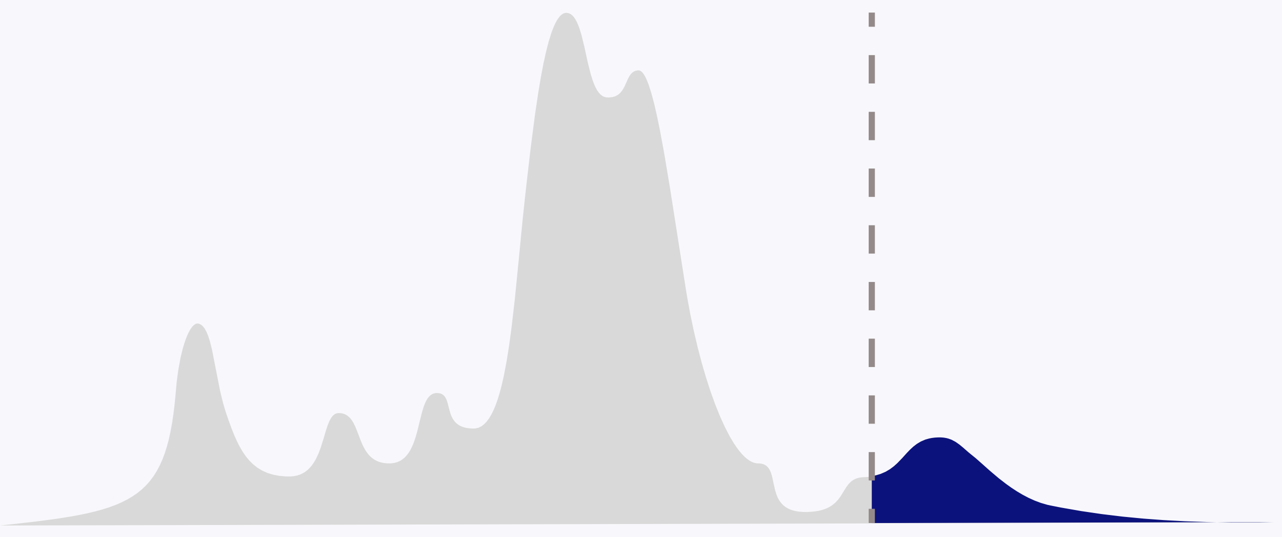

$\mathrm{CVaR}^{f}_{1-\beta}(p^{\pi}_{1})$

Part II

Data Debiasing and Hidden Mode Discovery

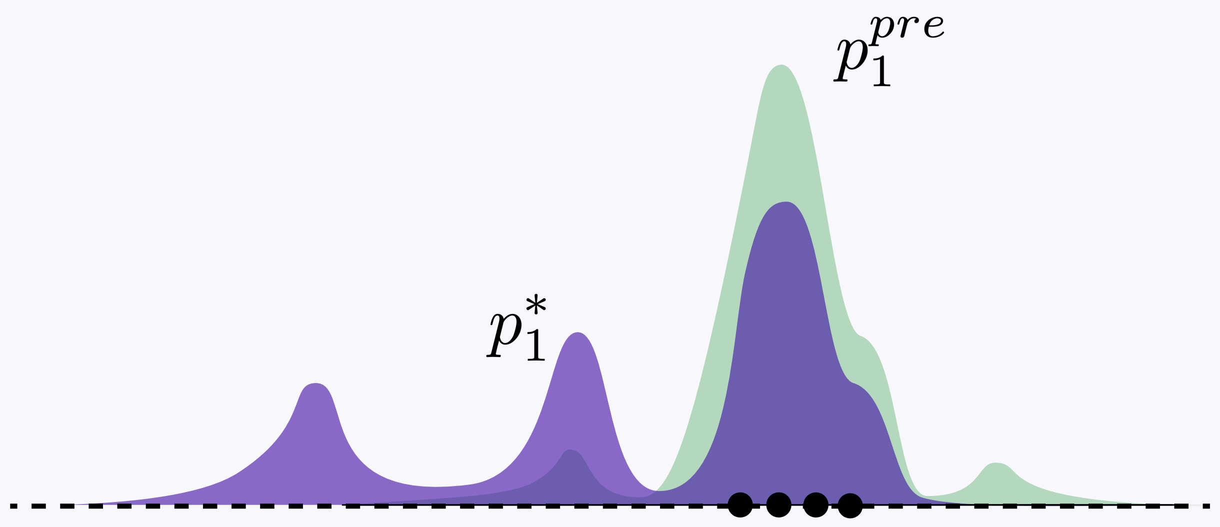

Part III

Out-of-Distribution Flow Modeling

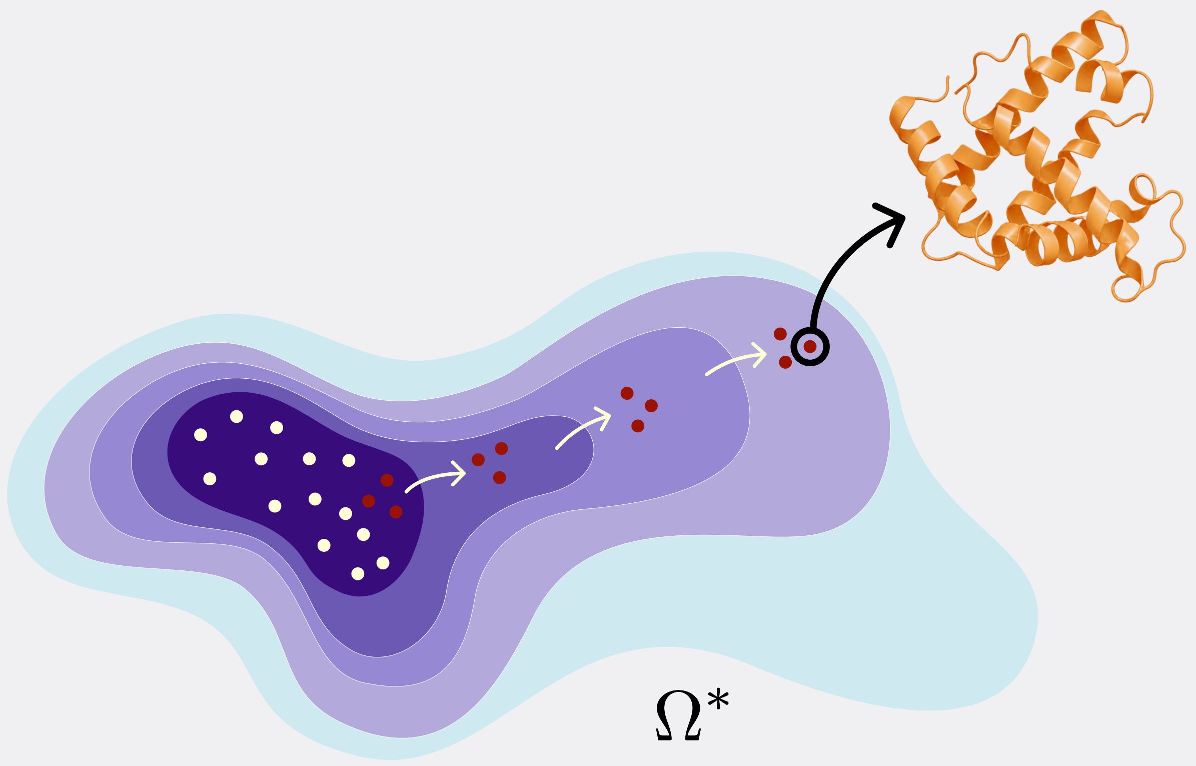

Constrained Entropic Flow Expansion: Key Reference

Verifier-Constrained Flow Expansion for Discovery Beyond the Data

[**Riccardo De Santi***, Kimon Protopapas*, Ya-Ping Hsieh, and Andreas Krause]

ICLR 2026

Mode Discovery via Regularized Entropic Expansion

Pre-trained generative model

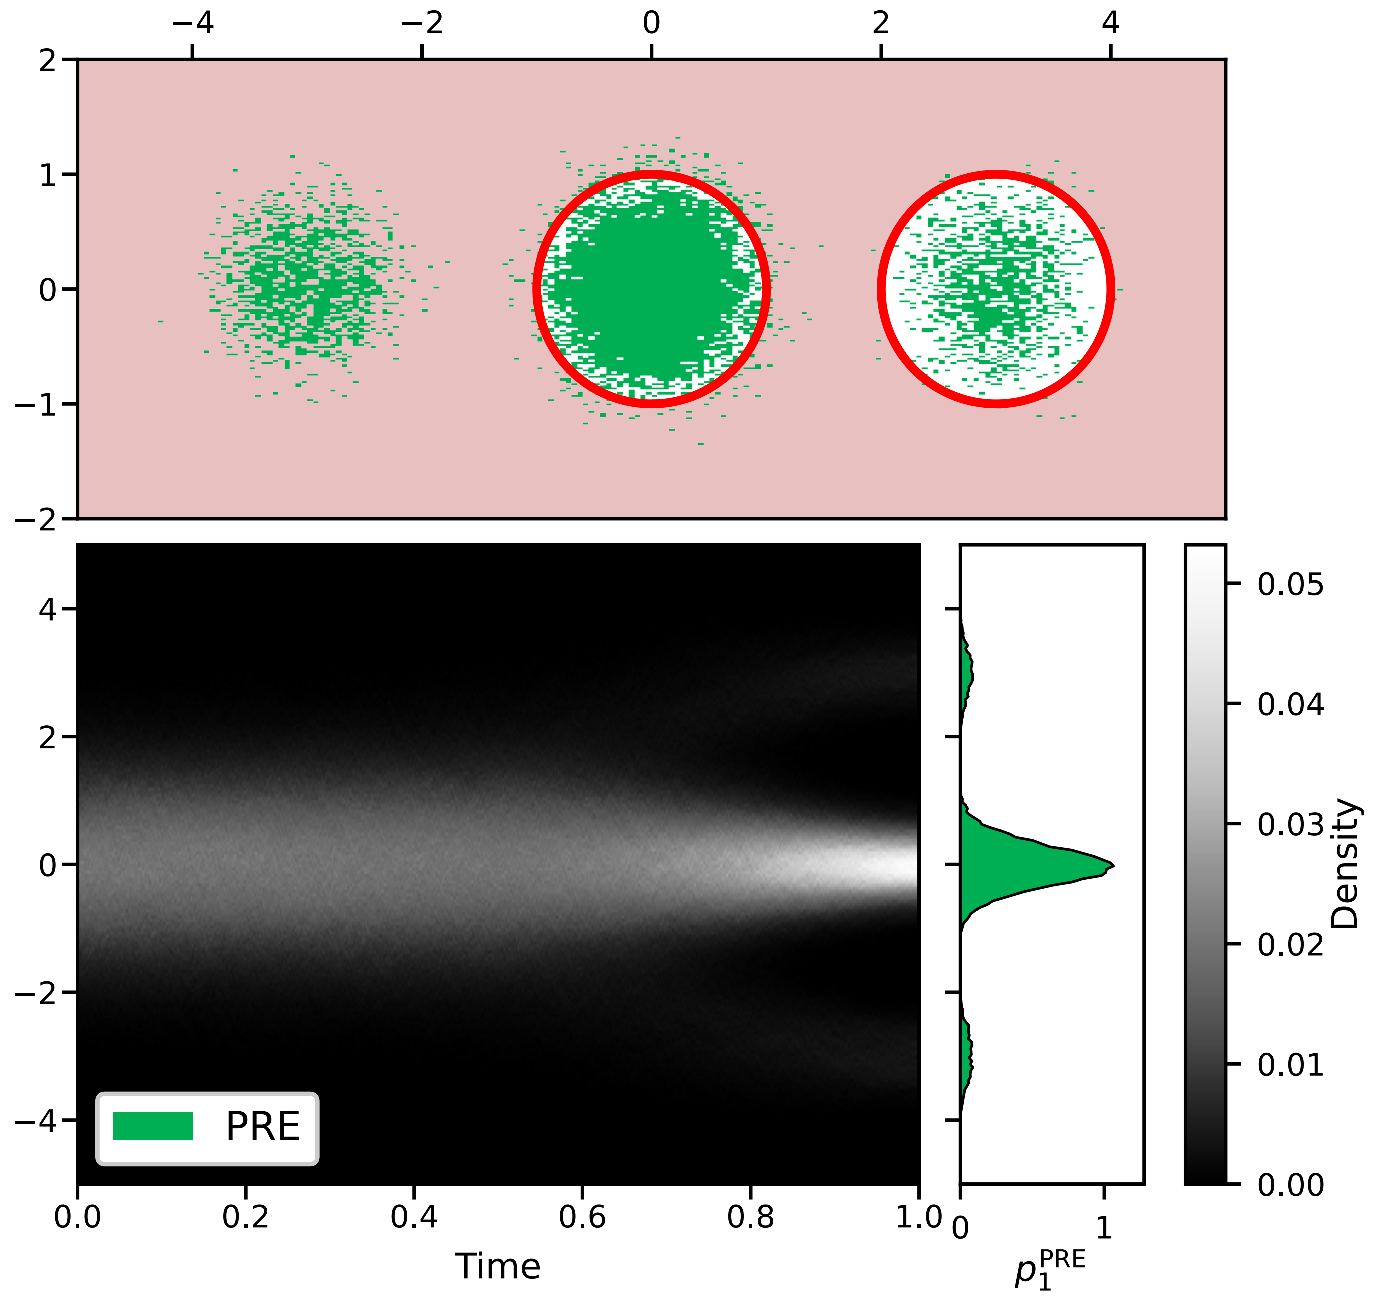

Mode Discovery via Regularized Entropic Expansion

Pre-trained generative model

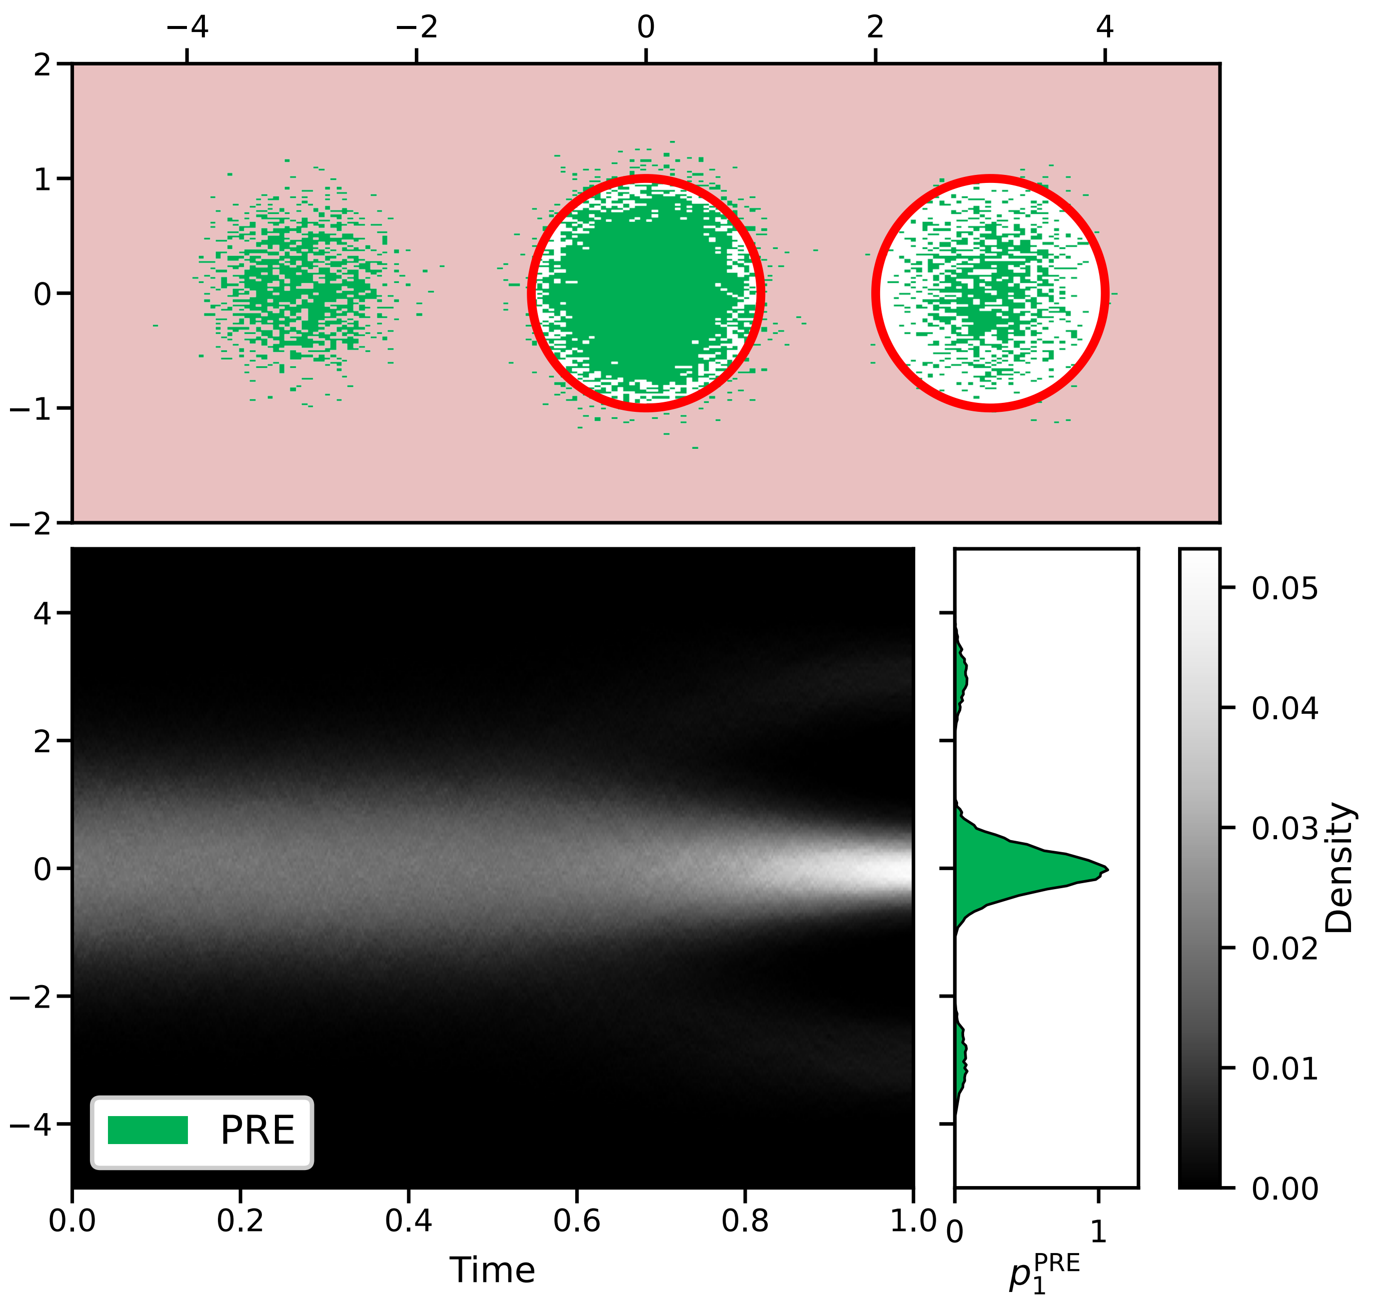

FDC

Expanded generative model

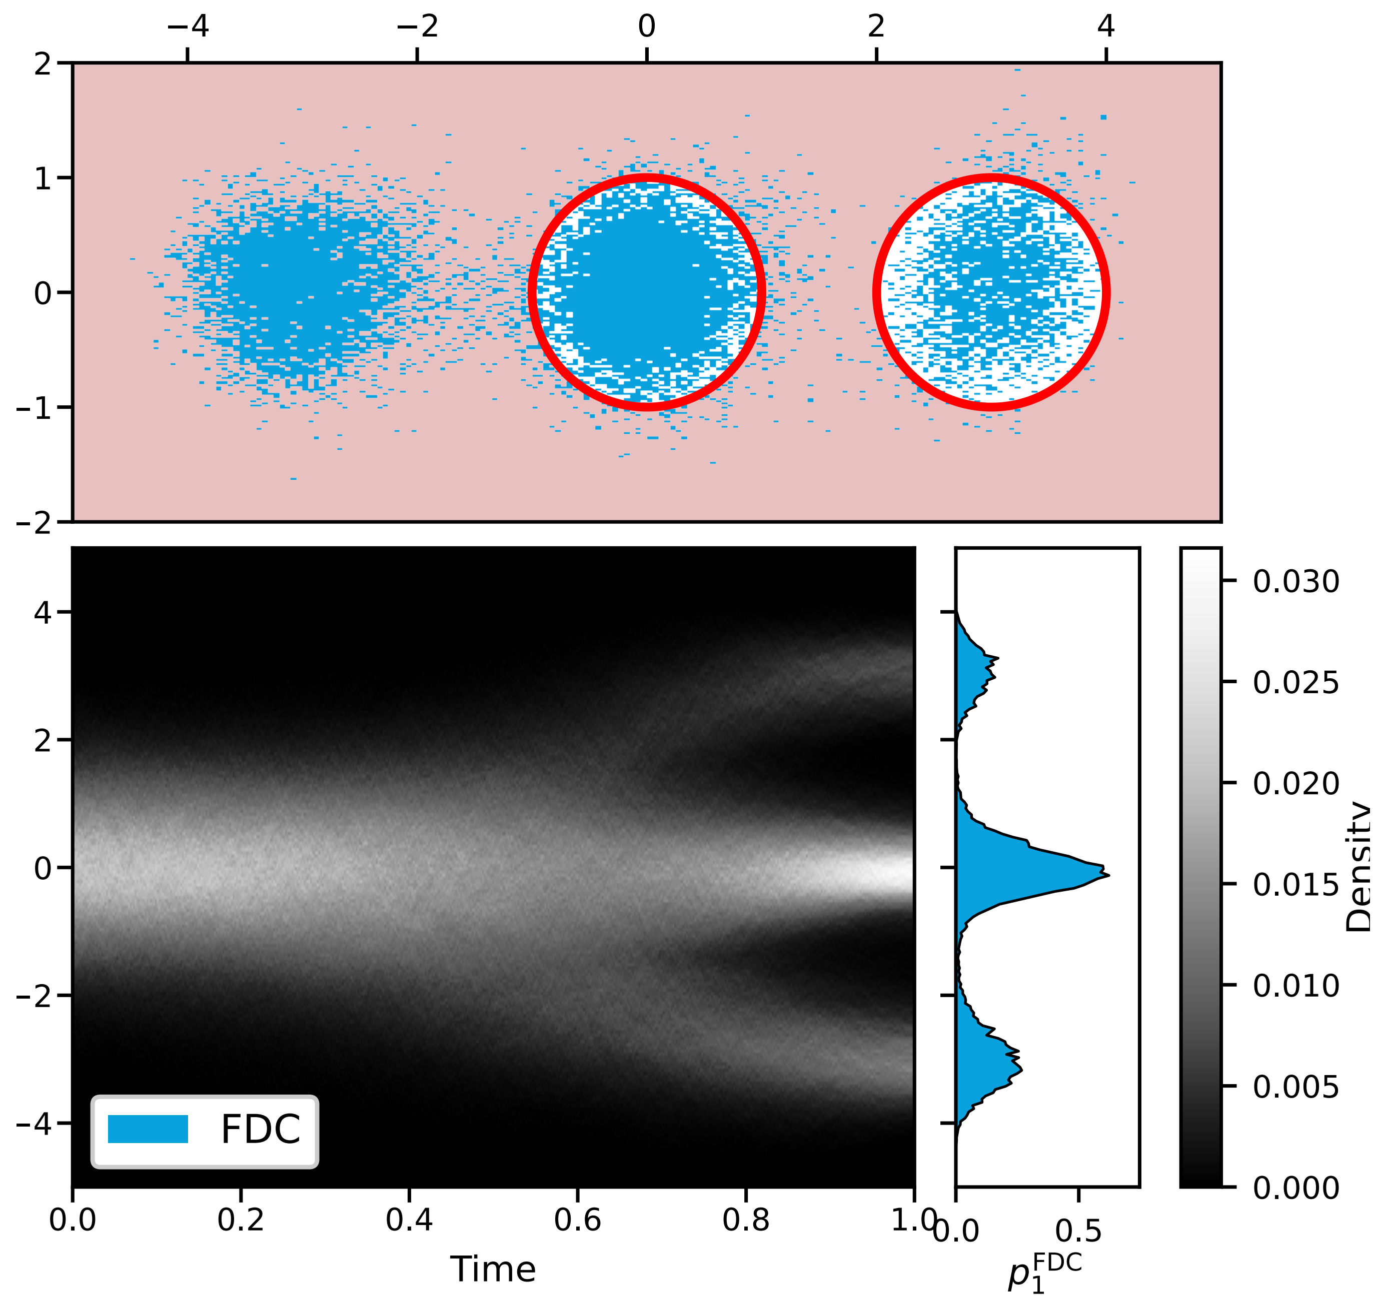

Mode Discovery via Regularized Entropic Expansion

Pre-trained generative model

Expanded generative model

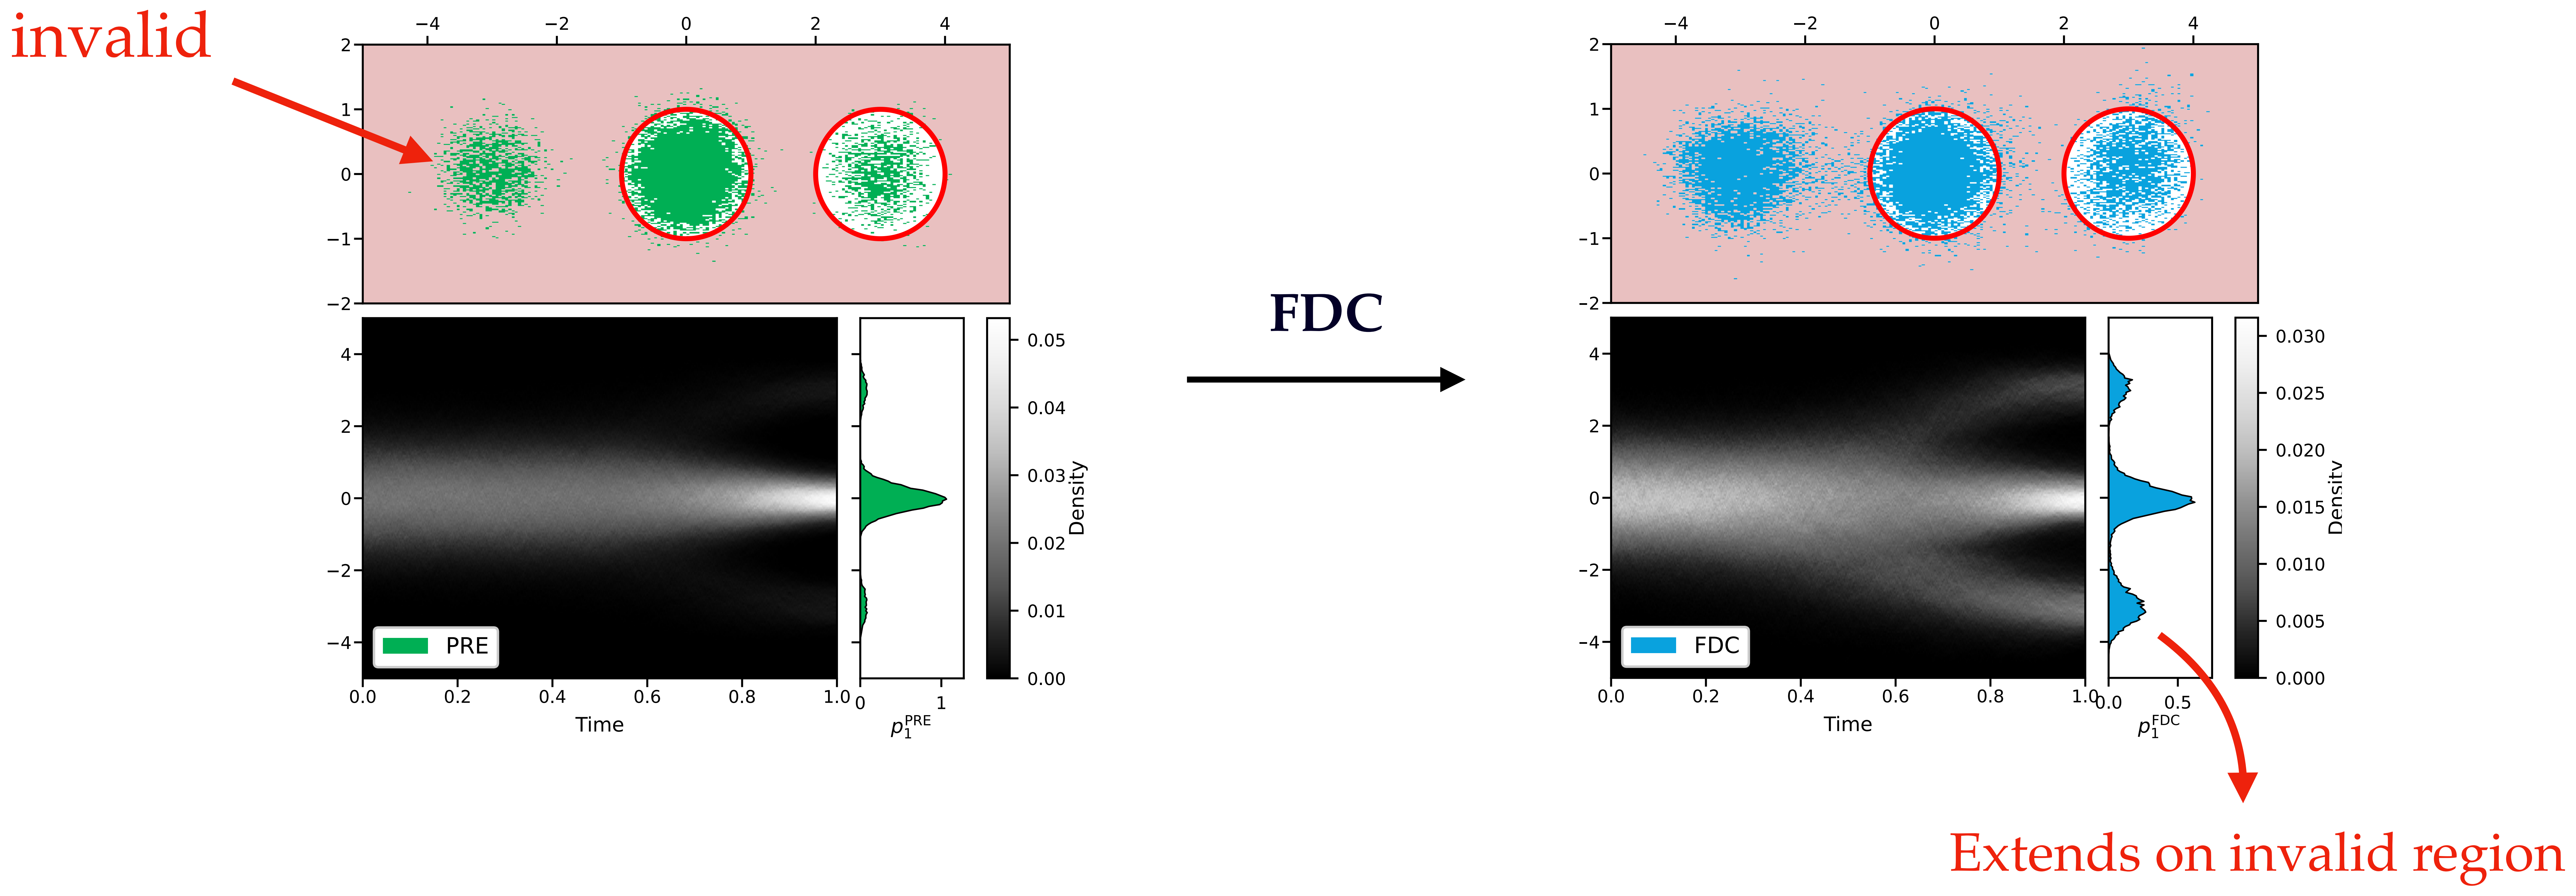

Stretching the models "creative power" via entropic expansion might increase density in invalid regions

Verifier-Constrained Flow Expansion

Fragmented molecules
(hallucinations)

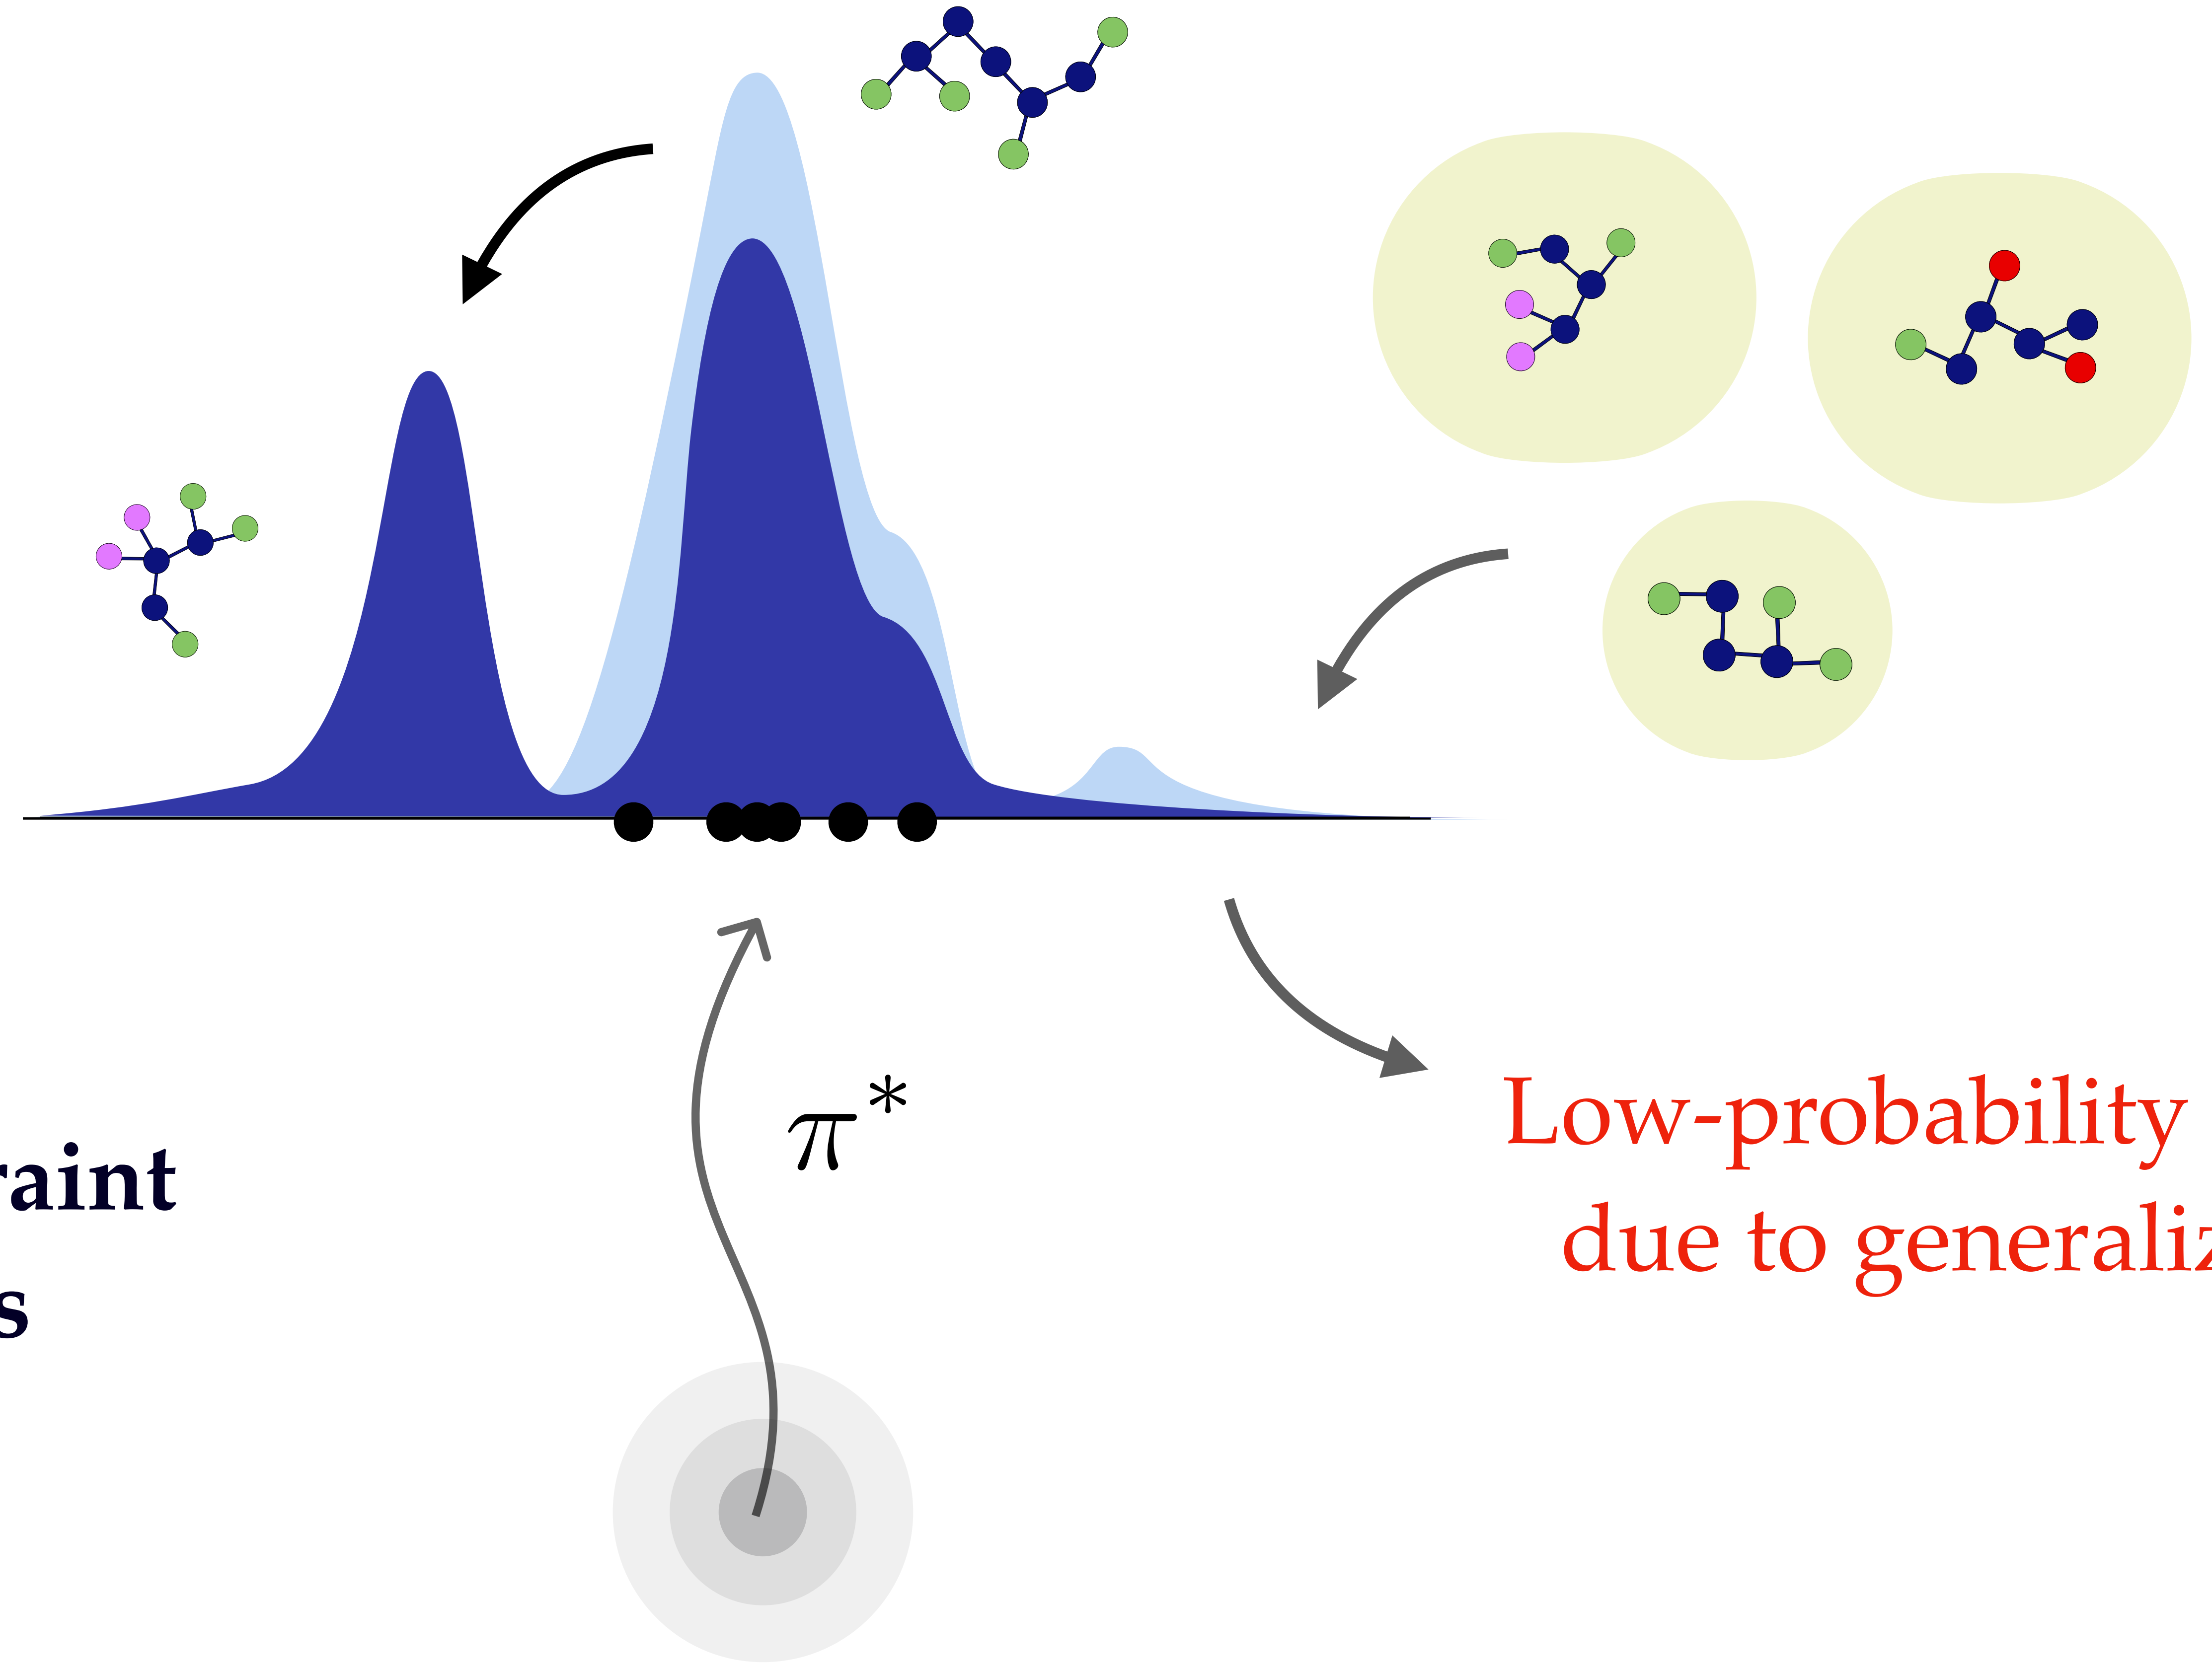

Low-probability modes can be
due to generalization errors

Key idea: use **verifiers to constraint expansion over valid regions**

Verifier

$$v : \mathcal{X} \rightarrow \{0, 1\}$$

Verifier-Constrained Flow Expansion

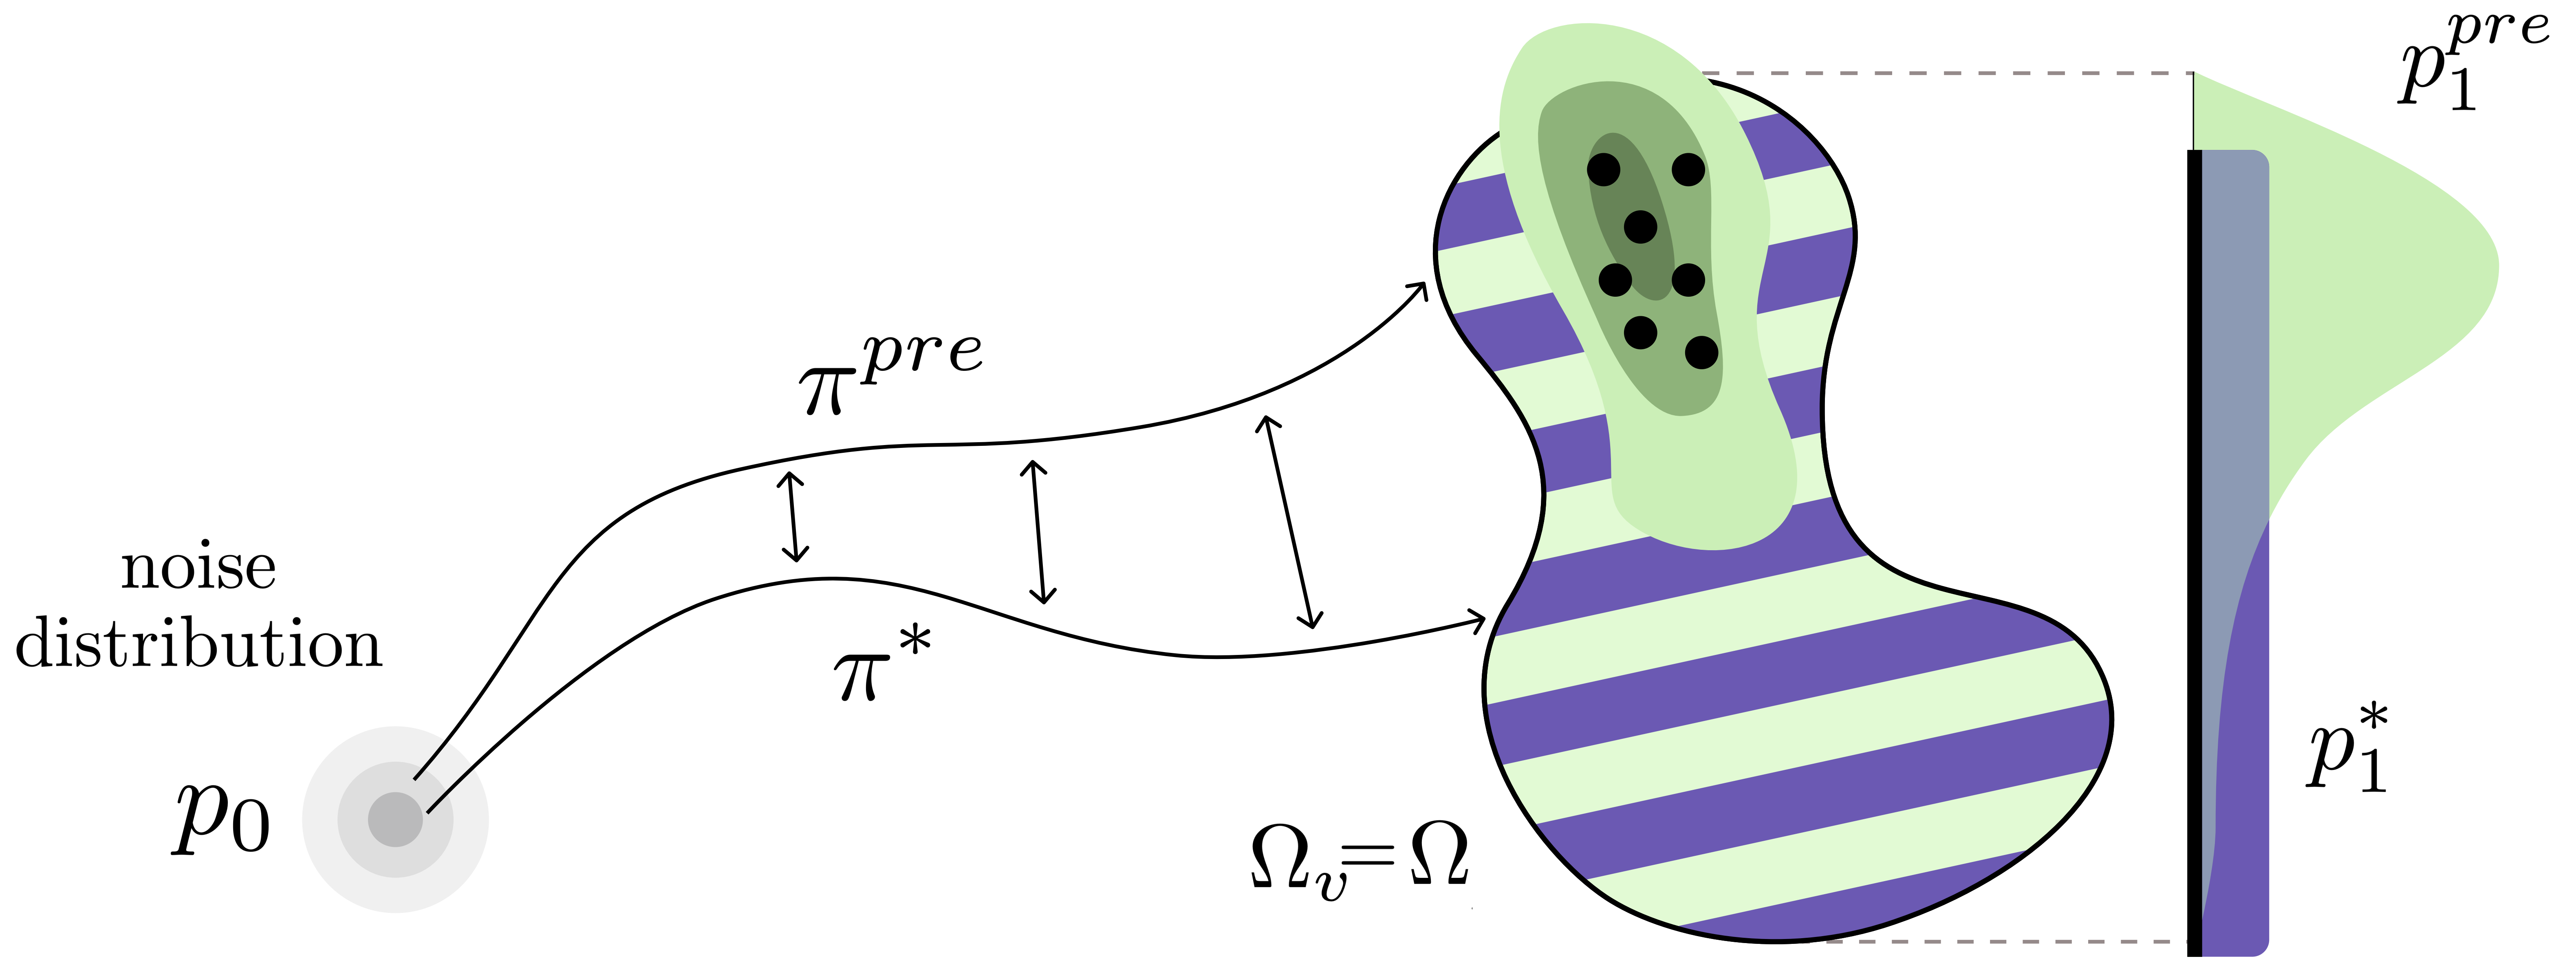

Verifier-Constrained Entropic Expansion

$$\pi^* \in \underset{\pi}{\operatorname{argmax}} \; \mathcal{H}(p_1^\pi) - \alpha \mathbf{KL}(p_1^\pi \,\|\, p_1^{pre}) \text{ s.t. } \mathbb{E}_{x \sim p_1^\pi}[v(x)] = 1$$

KL-reg. entropic expansion

verifier constraint

(Local) Verifier-Constrained Flow Expansion (L-FE)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \dots, K$:

Set $$\nabla f_k := \nabla \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}} = -(1+\alpha) s_1^{\pi_{k-1}} + s_1^{\pi_{\text{pre}}}$$

Expansion Step

Fine-tune π_{k-1} via standard reward-guided fine-tuning:

$$\tilde{\pi}_k \longleftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi} [f_k(x)] - \eta_k \mathbf{KL}(p_1^\pi \,\|\, p_1^{k-1})$$

Return $\pi := \pi_K$

continued pre-training to obtain a **better generative prior for downstream discovery tasks**

(Local) Verifier-Constrained Flow Expansion (L-FE)

Init: $\pi_0 := \pi^{pre}$

For $k = 1, \dots, K$:

Set $$\nabla f_k := \nabla \frac{\delta \mathcal{G}(p_1^{\pi_{k-1}})}{\delta p_1^{\pi_{k-1}}} = -(1+\alpha) s_1^{\pi_{k-1}} + s_1^{\pi_{\text{pre}}}$$

Expansion Step

Fine-tune π_{k-1} via standard reward-guided fine-tuning:

$$\tilde{\pi}_k \leftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi} [f_k(x)] - \eta_k \mathbf{KL}(p_1^\pi \,\|\, p_1^{k-1})$$

Fine-tune $\tilde{\pi}_k$ to enforce verifier constraints:

$$\tilde{\pi}_k \leftarrow \underset{\pi}{\text{argmax}} \quad \mathbb{E}_{x \sim p_1^\pi} [\log v(x)] - \eta_k \mathbf{KL}(p_1^\pi \,\|\, p_1^{k-1})$$

Projection Step

Return $\pi := \pi_K$

continued pre-training to obtain a **better generative prior for downstream discovery tasks**

Mode Discovery via Constrained Entropic Expansion

Pre-trained generative model

invalid

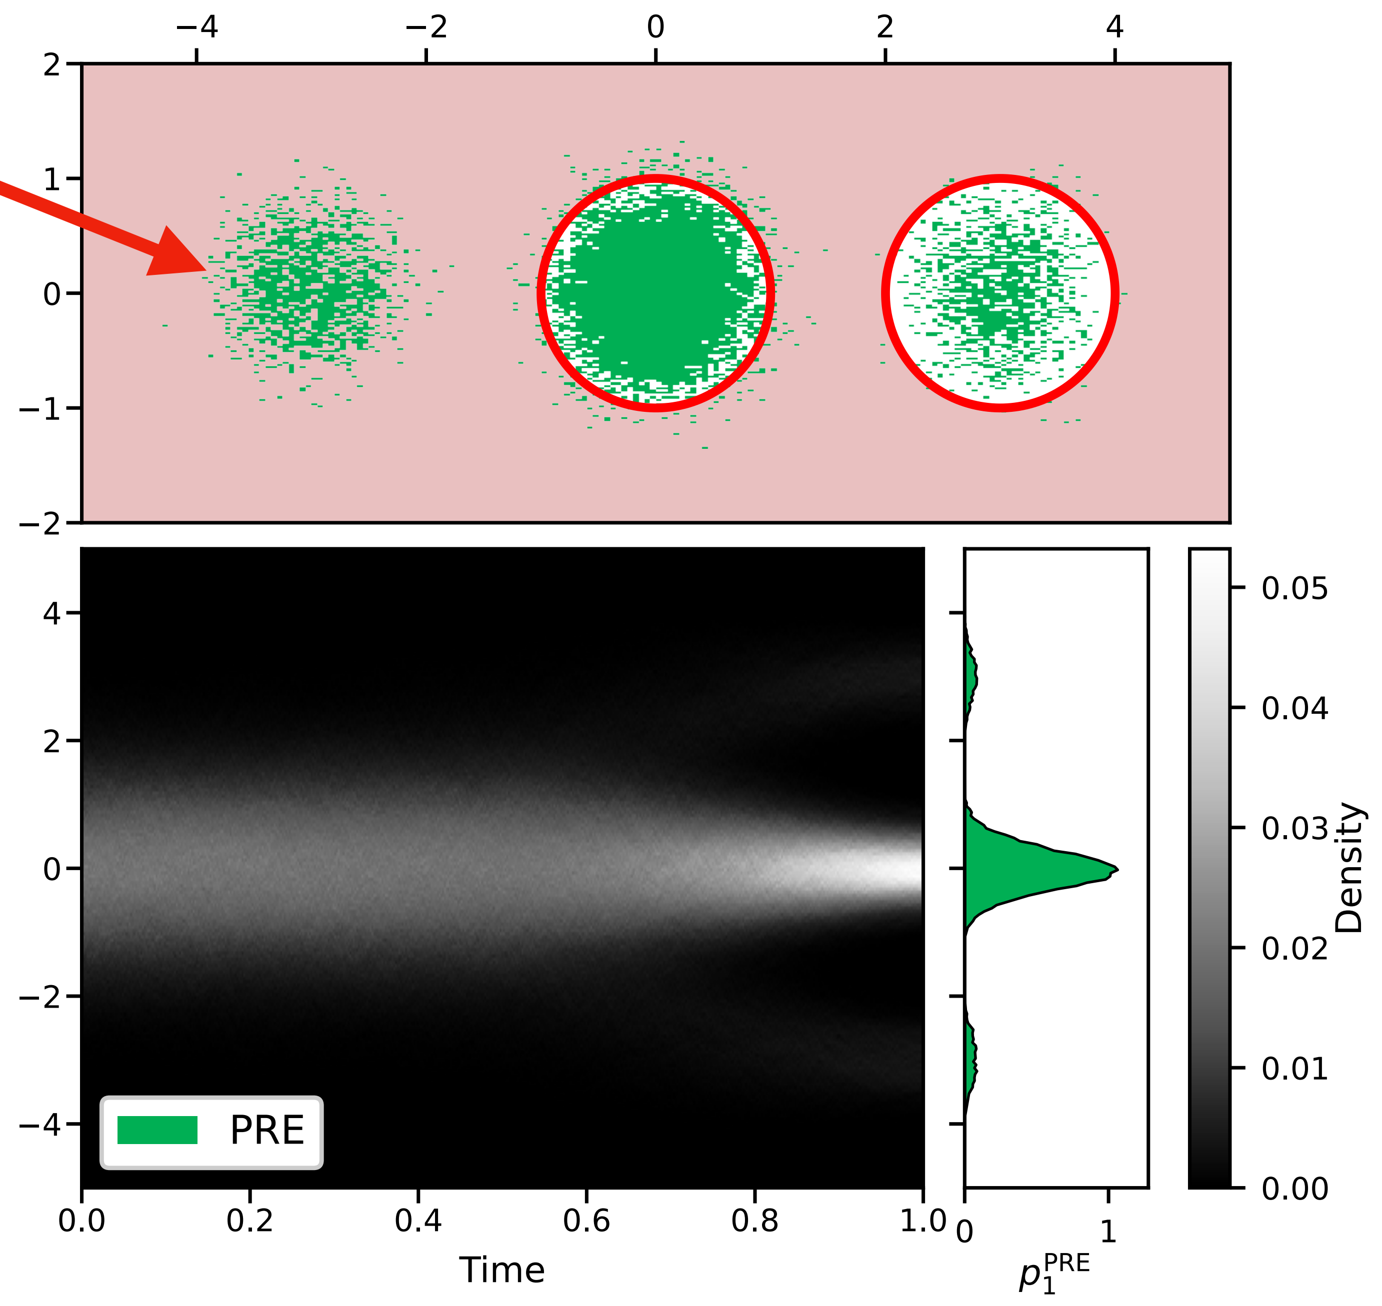

L-FE

Verifier-Constrained Expansion

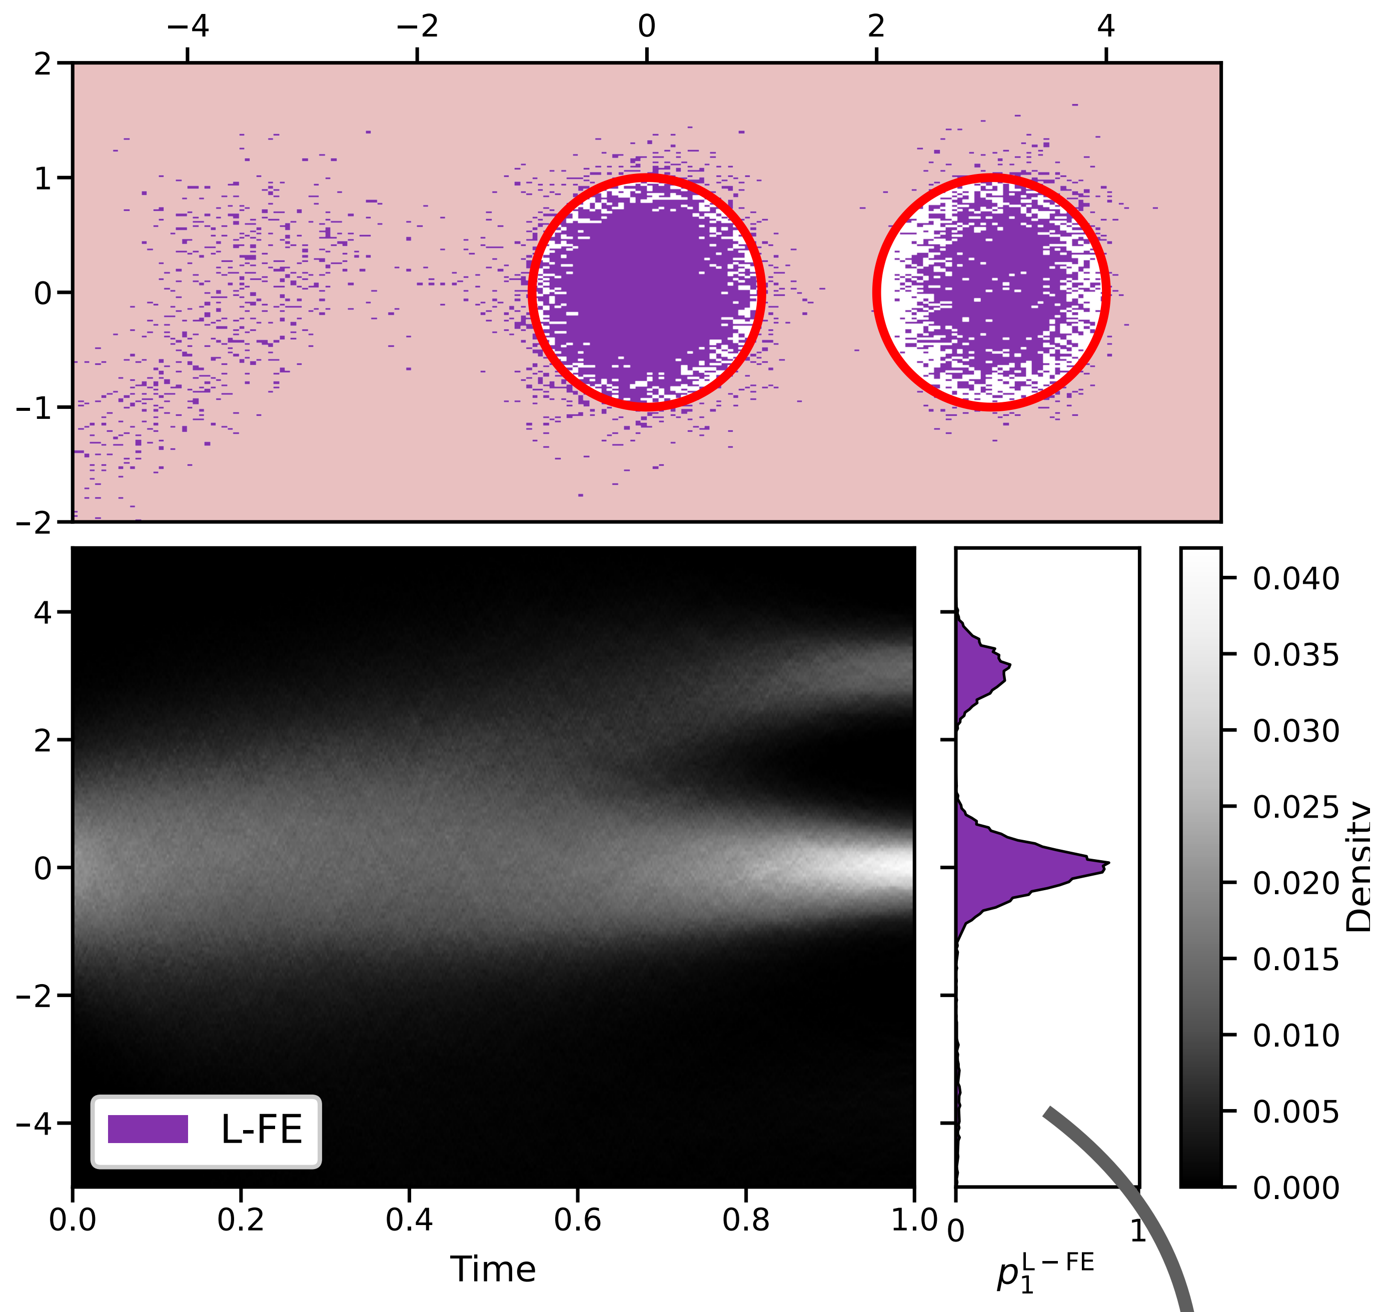

L-FE prevents expansion on invalid region

Flow Expansion for De Novo Molecular Design

(weak) verifier = interatomic distance bound

validity metric = RDKit Sanitize + interatomic distance bound

GEOM-Drugs

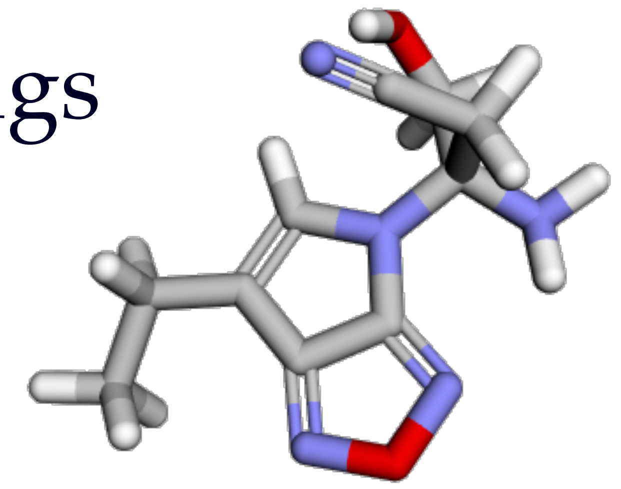

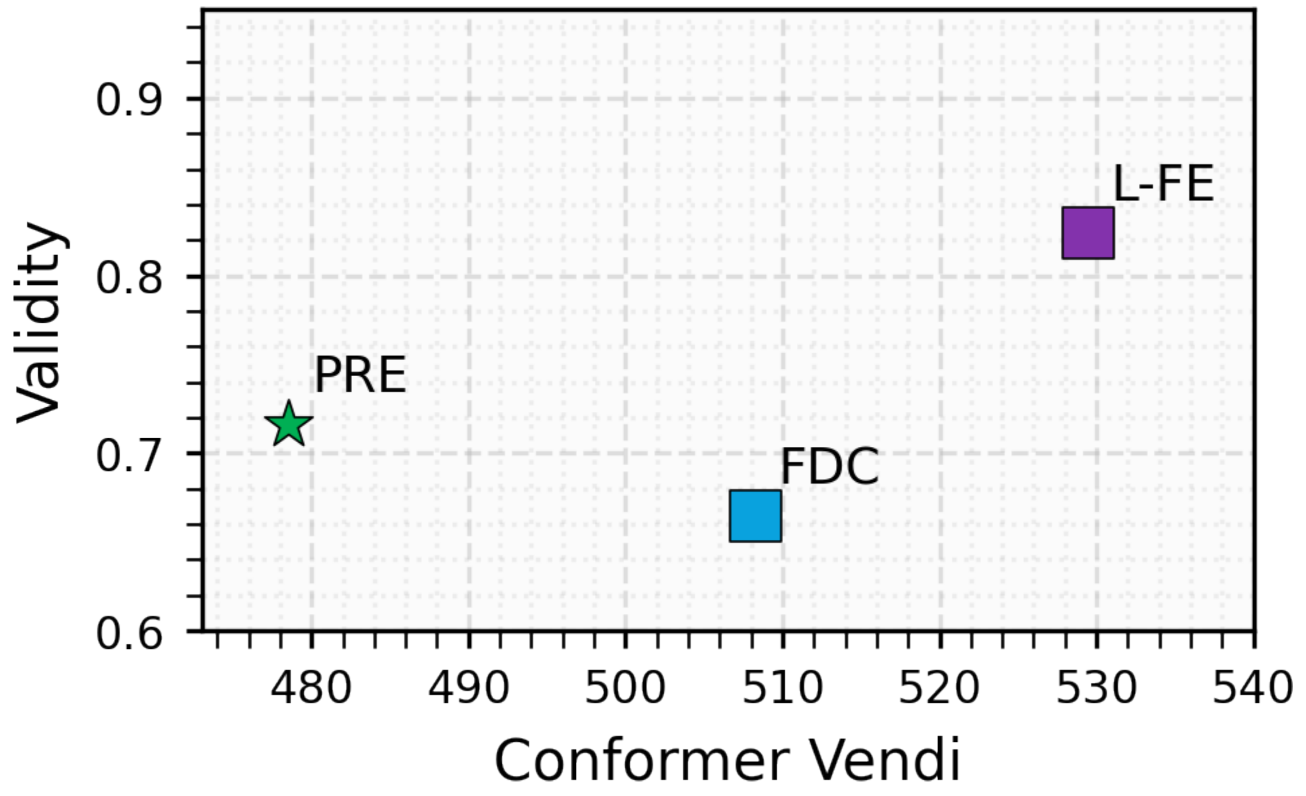

	Diversity	Validity (%)
PRE	476	72
FDC	508	66
L-FE	**529**	**82**

Flow Expansion for De Novo Molecular Design

(weak) verifier = interatomic distance bound

validity metric = RDKit Sanitize + interatomic distance bound

GEOM-Drugs

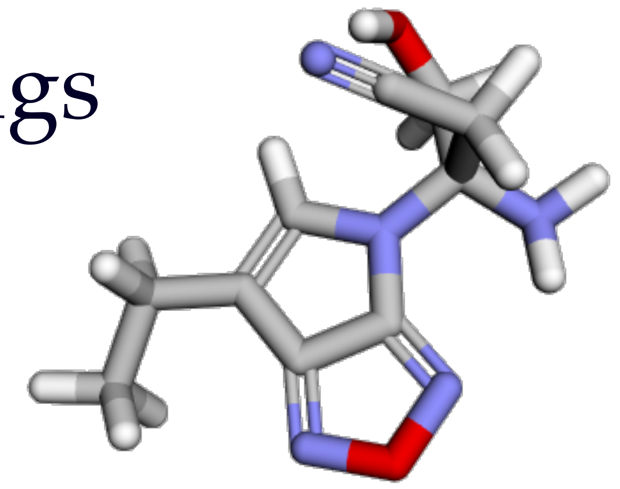

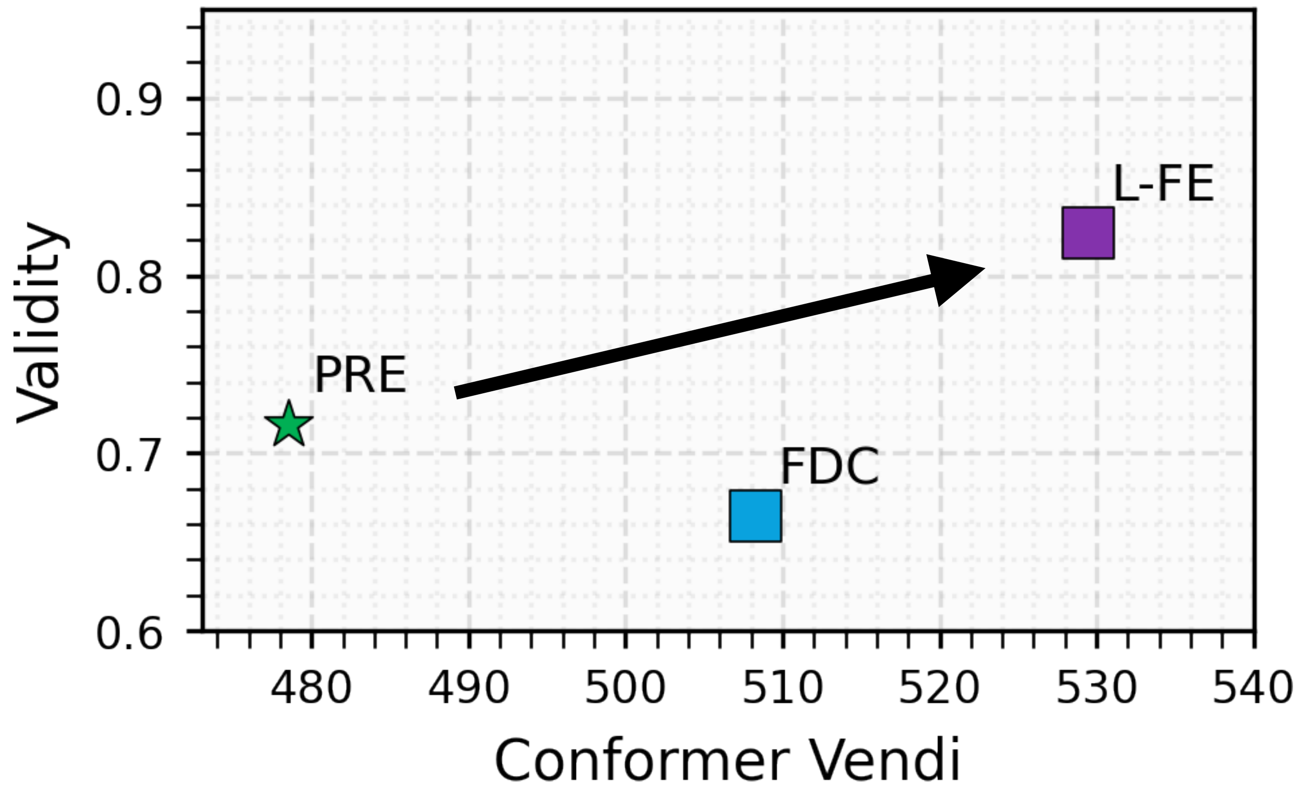

	Diversity	Validity (%)
PRE	476	72
FDC	508	66
L-FE	**529**	**82**

Takeaway

Flow Expansion can **increases diversity and validity** of pre-trained generative models

Out-of-Distribution Flow Modeling: Key Reference

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]

Preprint.

Diffusion and Flow Modeling: Standard Pre-training

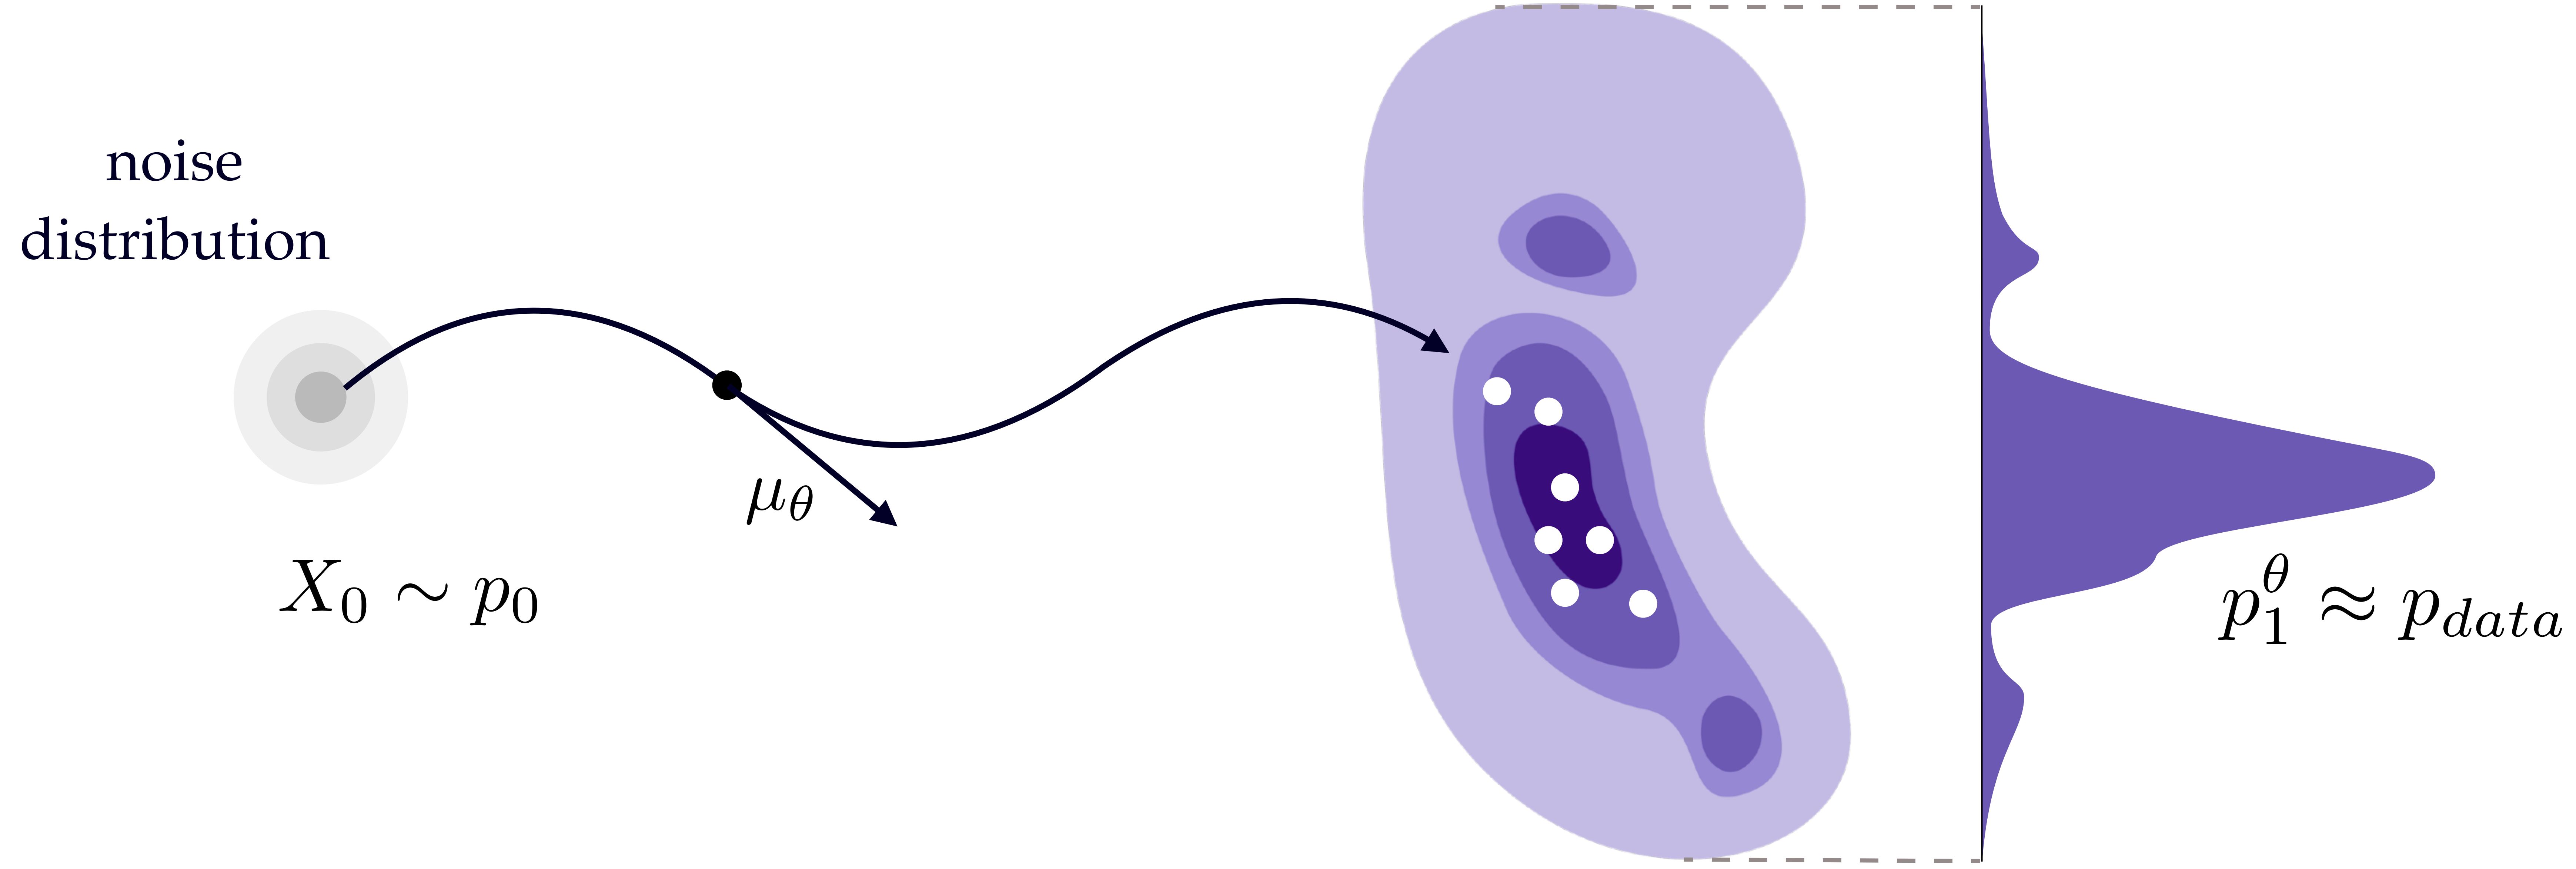

Distribution Matching: Standard Flow Pre-training

learn θ s.t. $p_1^\theta \approx p_{data}$

↓

Implement via e.g., Flow/Score Matching

Limitation of Standard Pre-training for Discovery

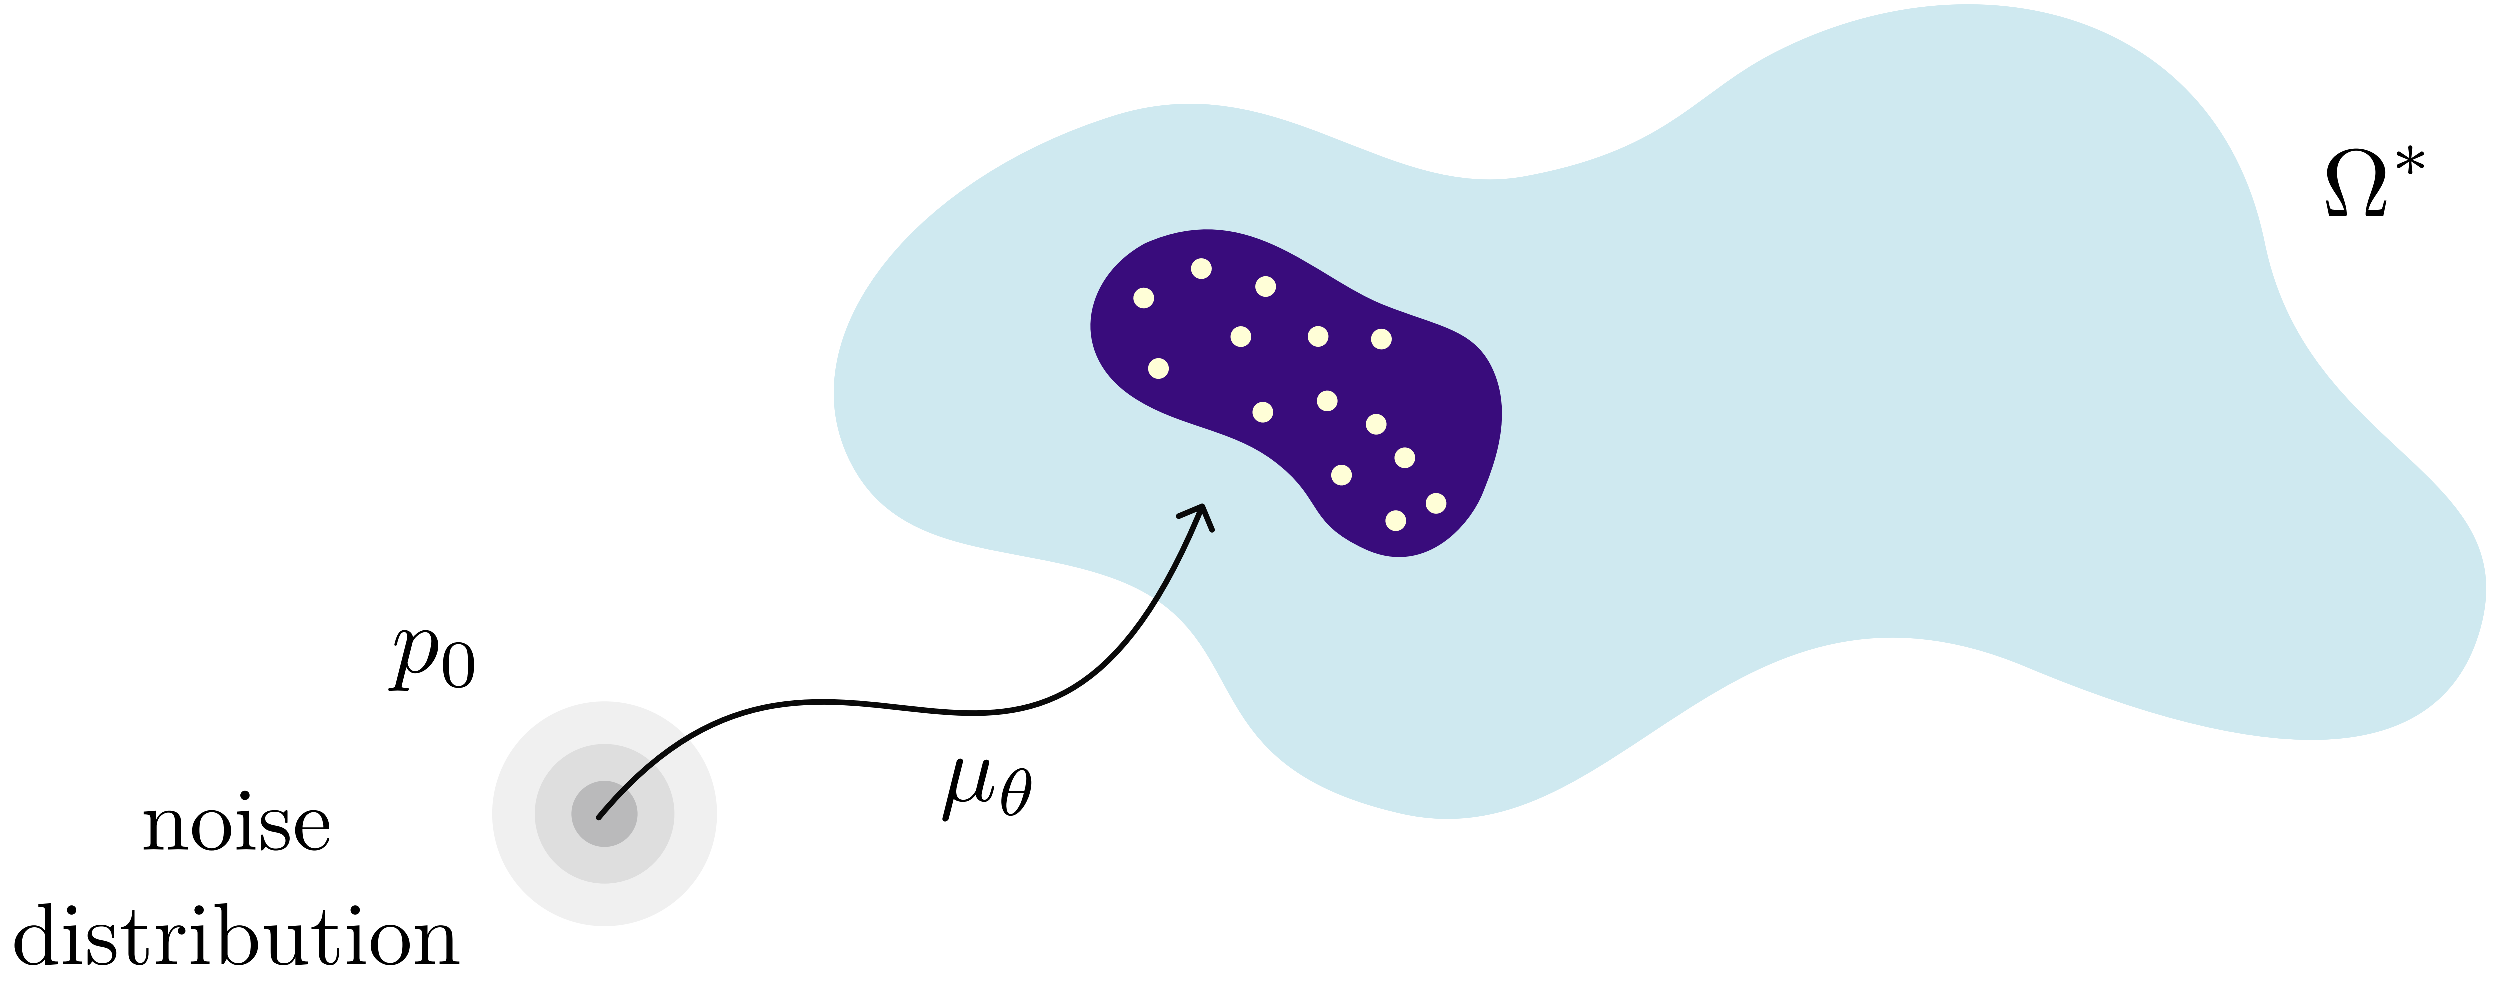

Distribution Matching: Standard Flow Pre-training

learn θ s.t. $p_1^\theta \approx p_{data}$

Limitation of Standard Pre-training for Discovery

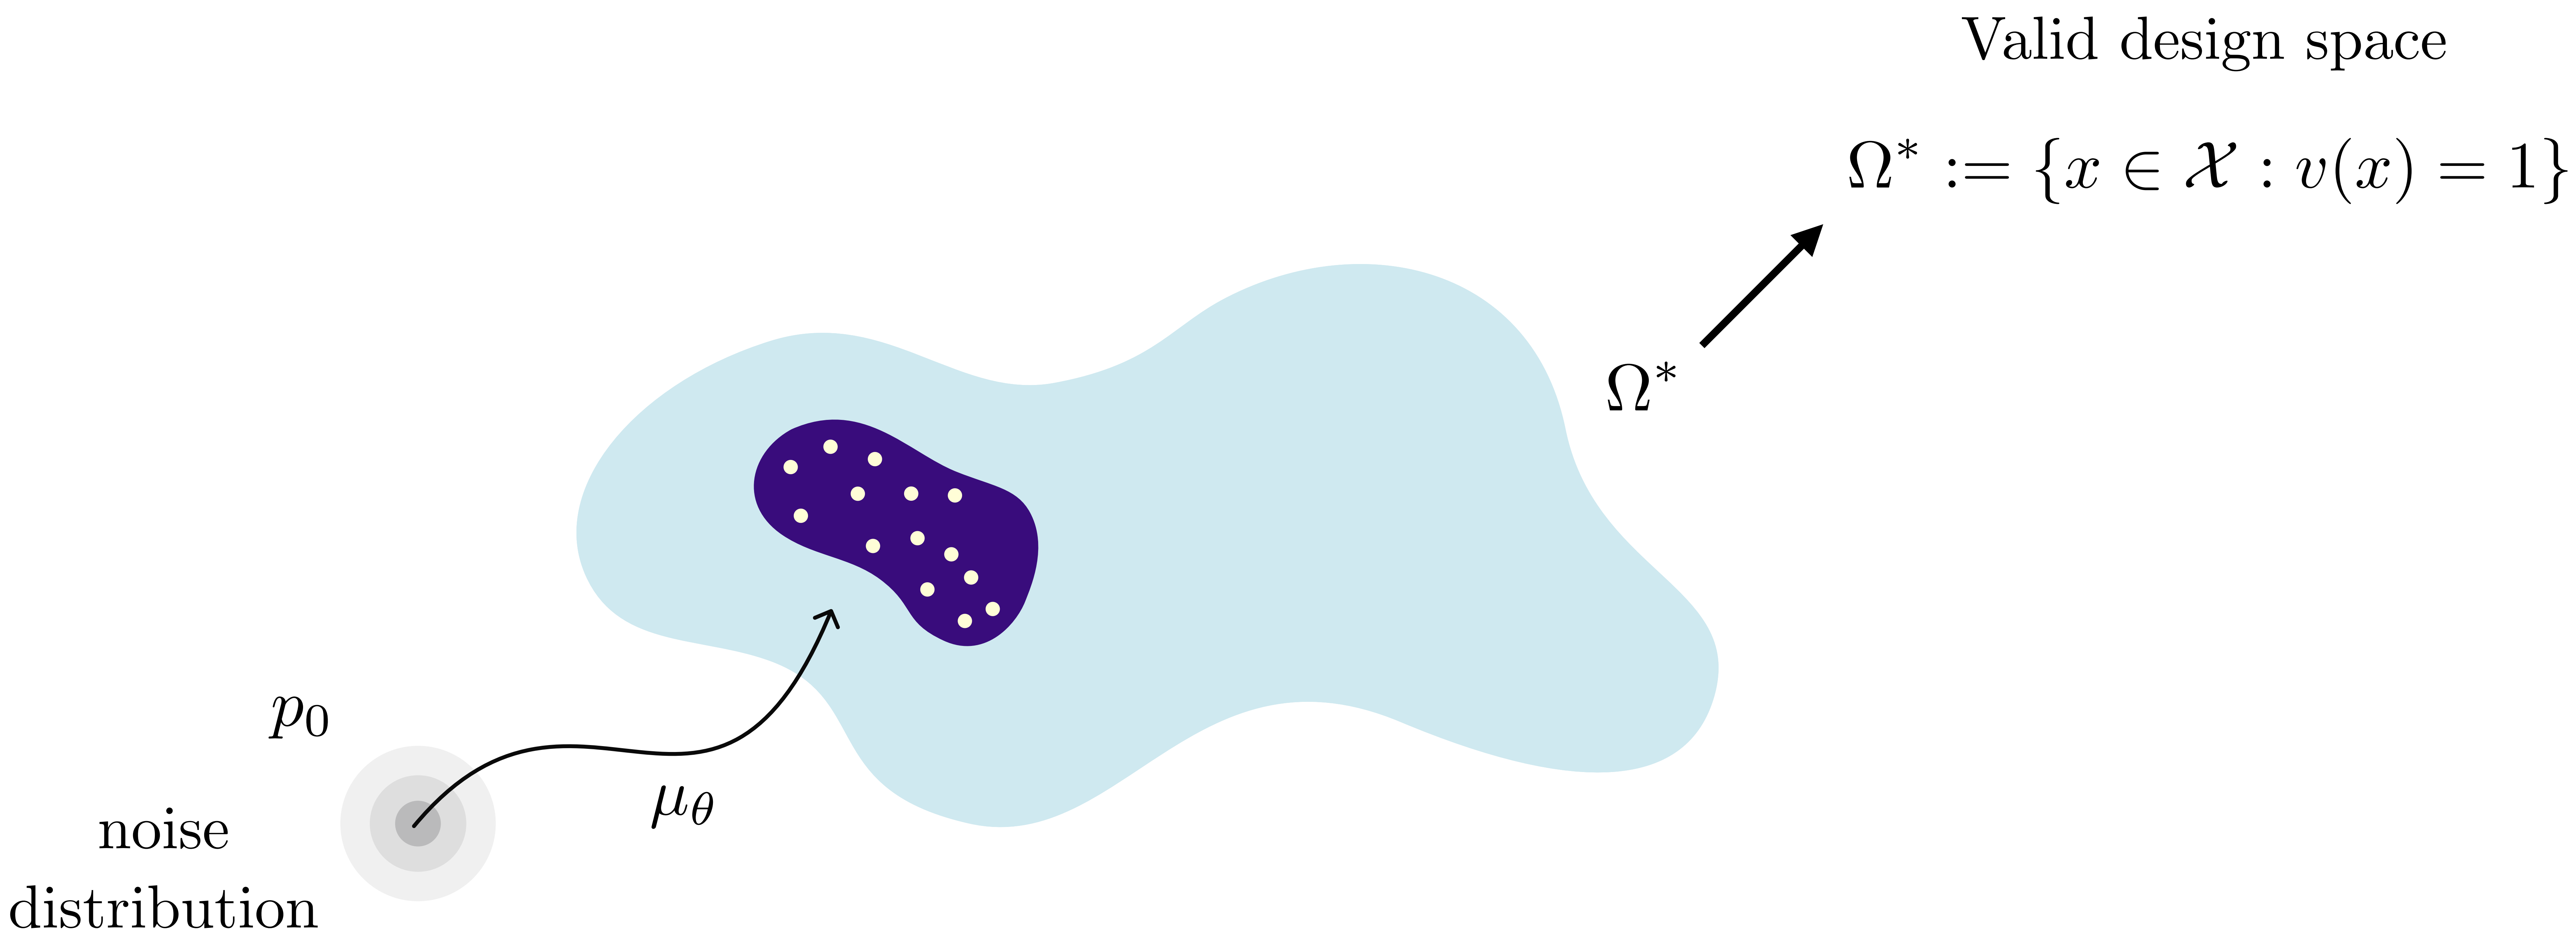

Distribution Matching: Standard Flow Pre-training

learn θ s.t. $p_1^\theta \approx p_{data}$

Limitation of Standard Pre-training for Discovery

How to formalize poor coverage of generative prior?

Valid design space

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

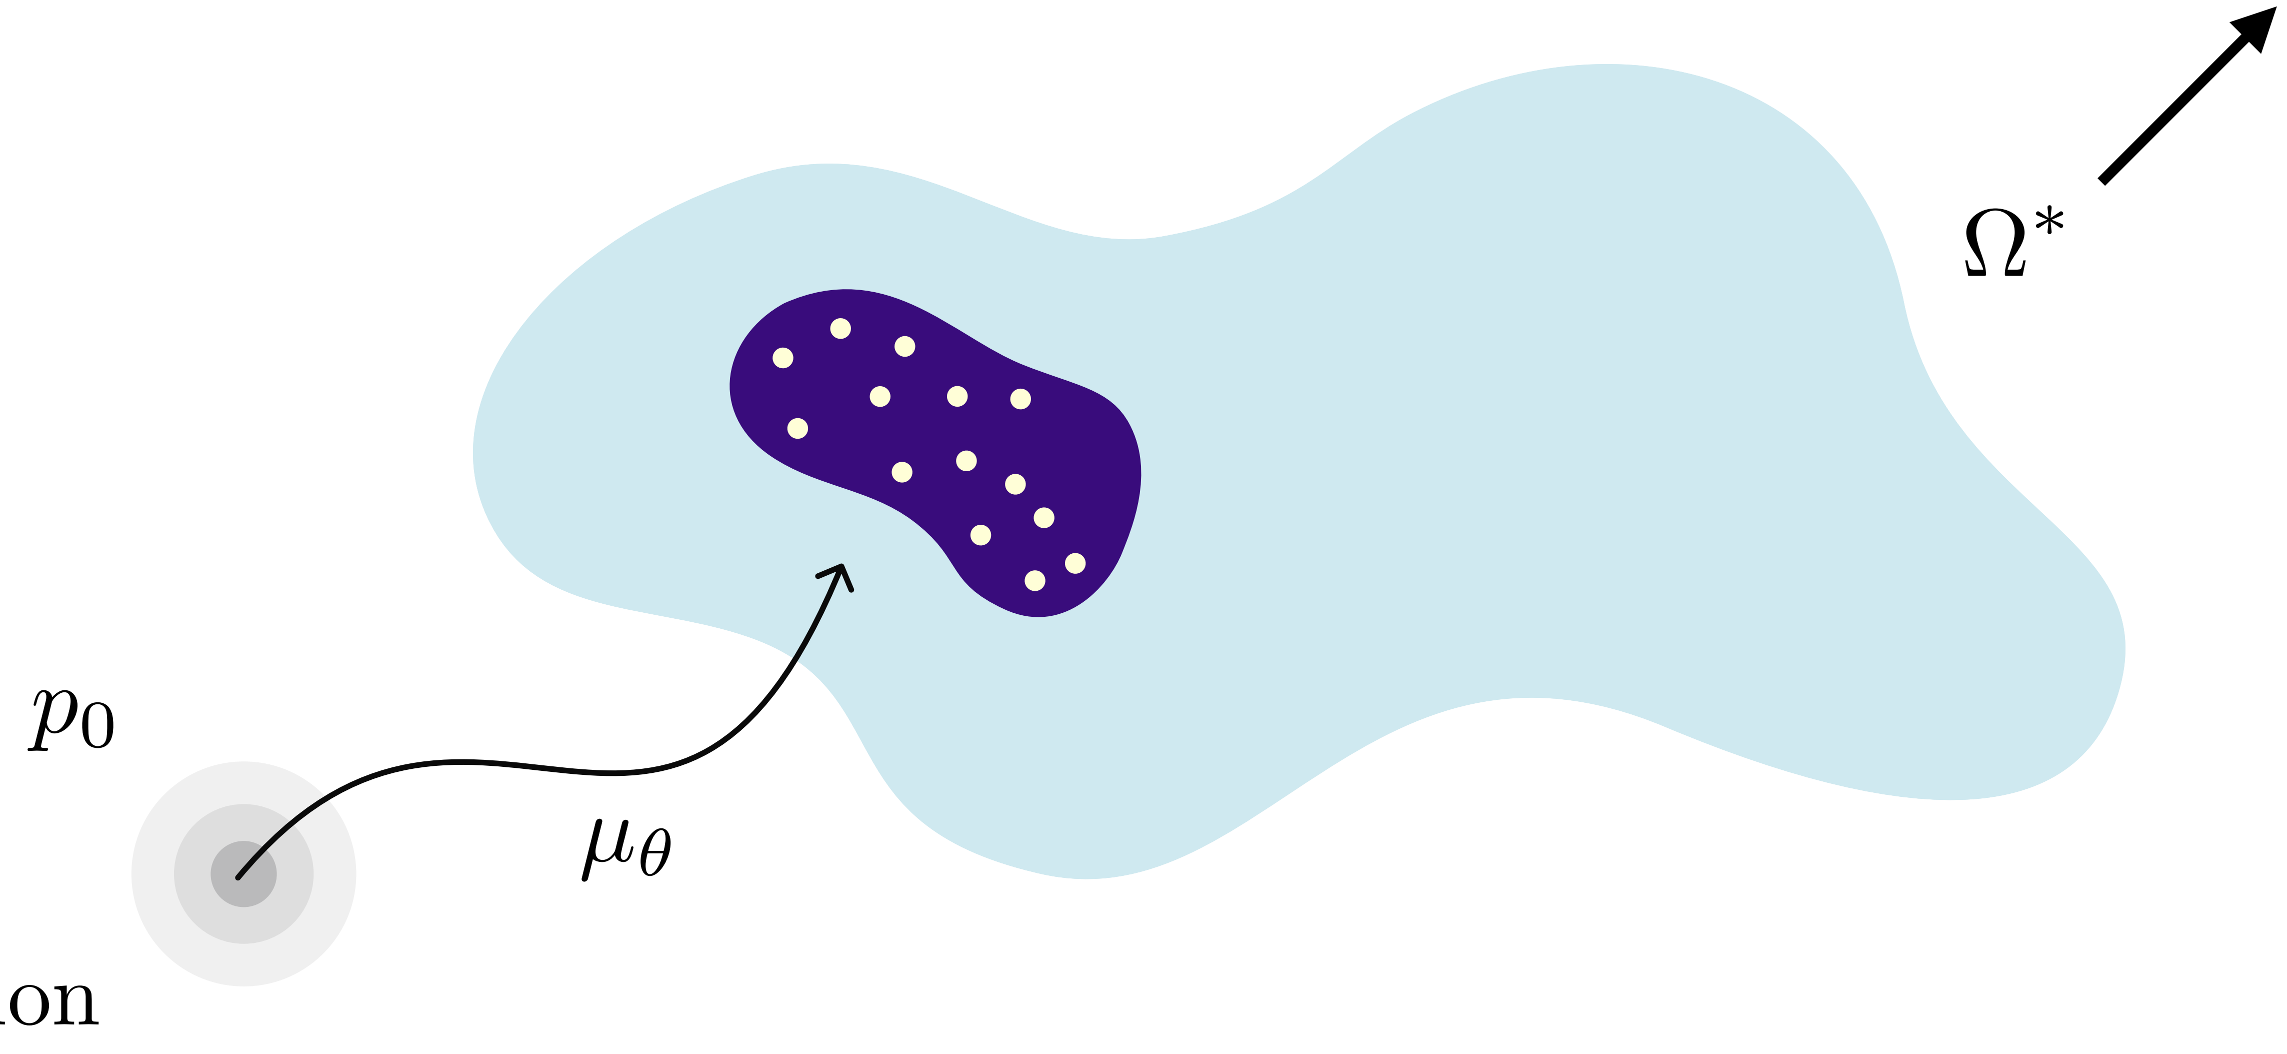

Distribution Matching: Standard Flow Pre-training

learn θ s.t. $p_1^\theta \approx p_{data}$

Generable Set

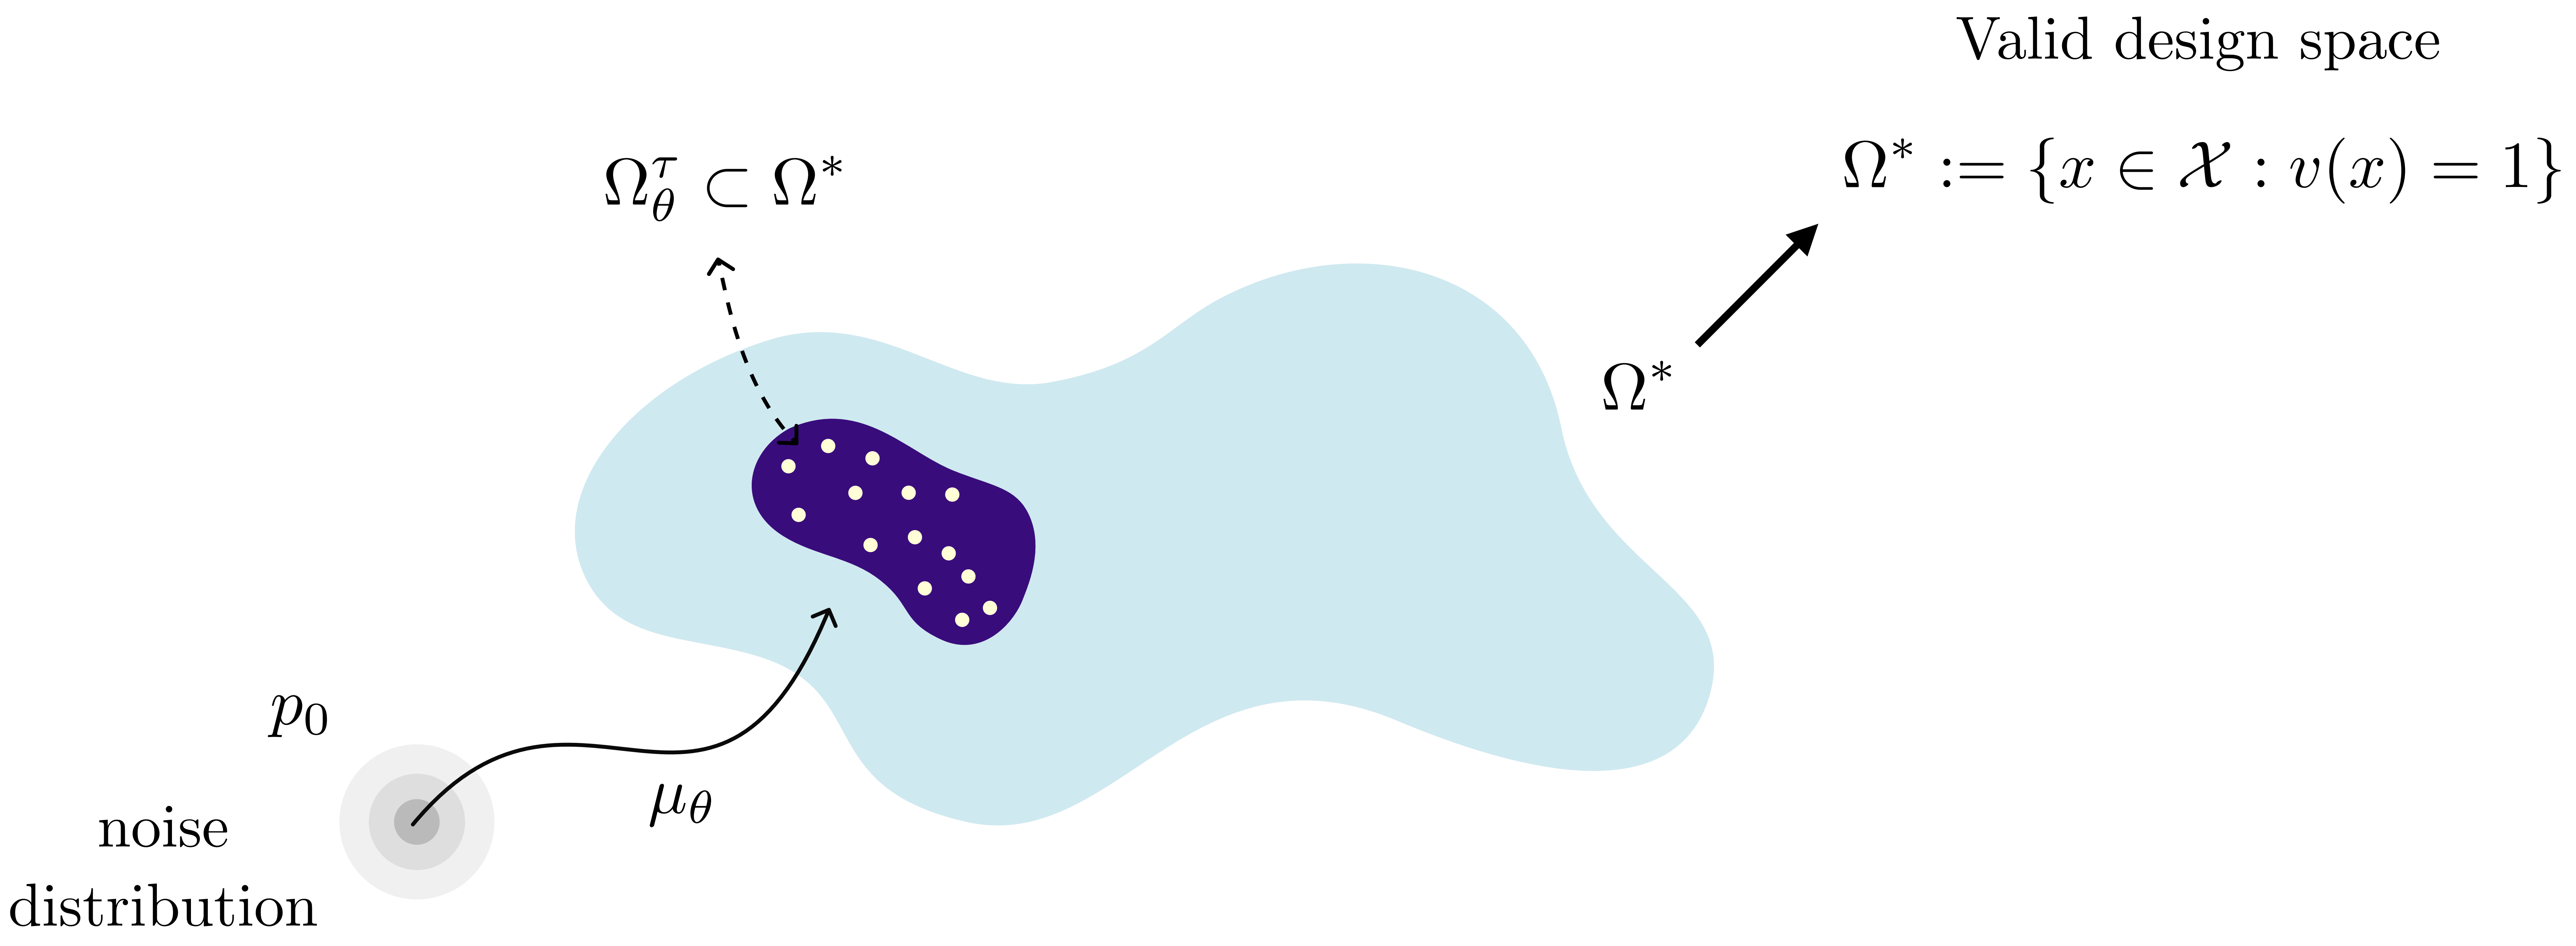

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

Generable Set

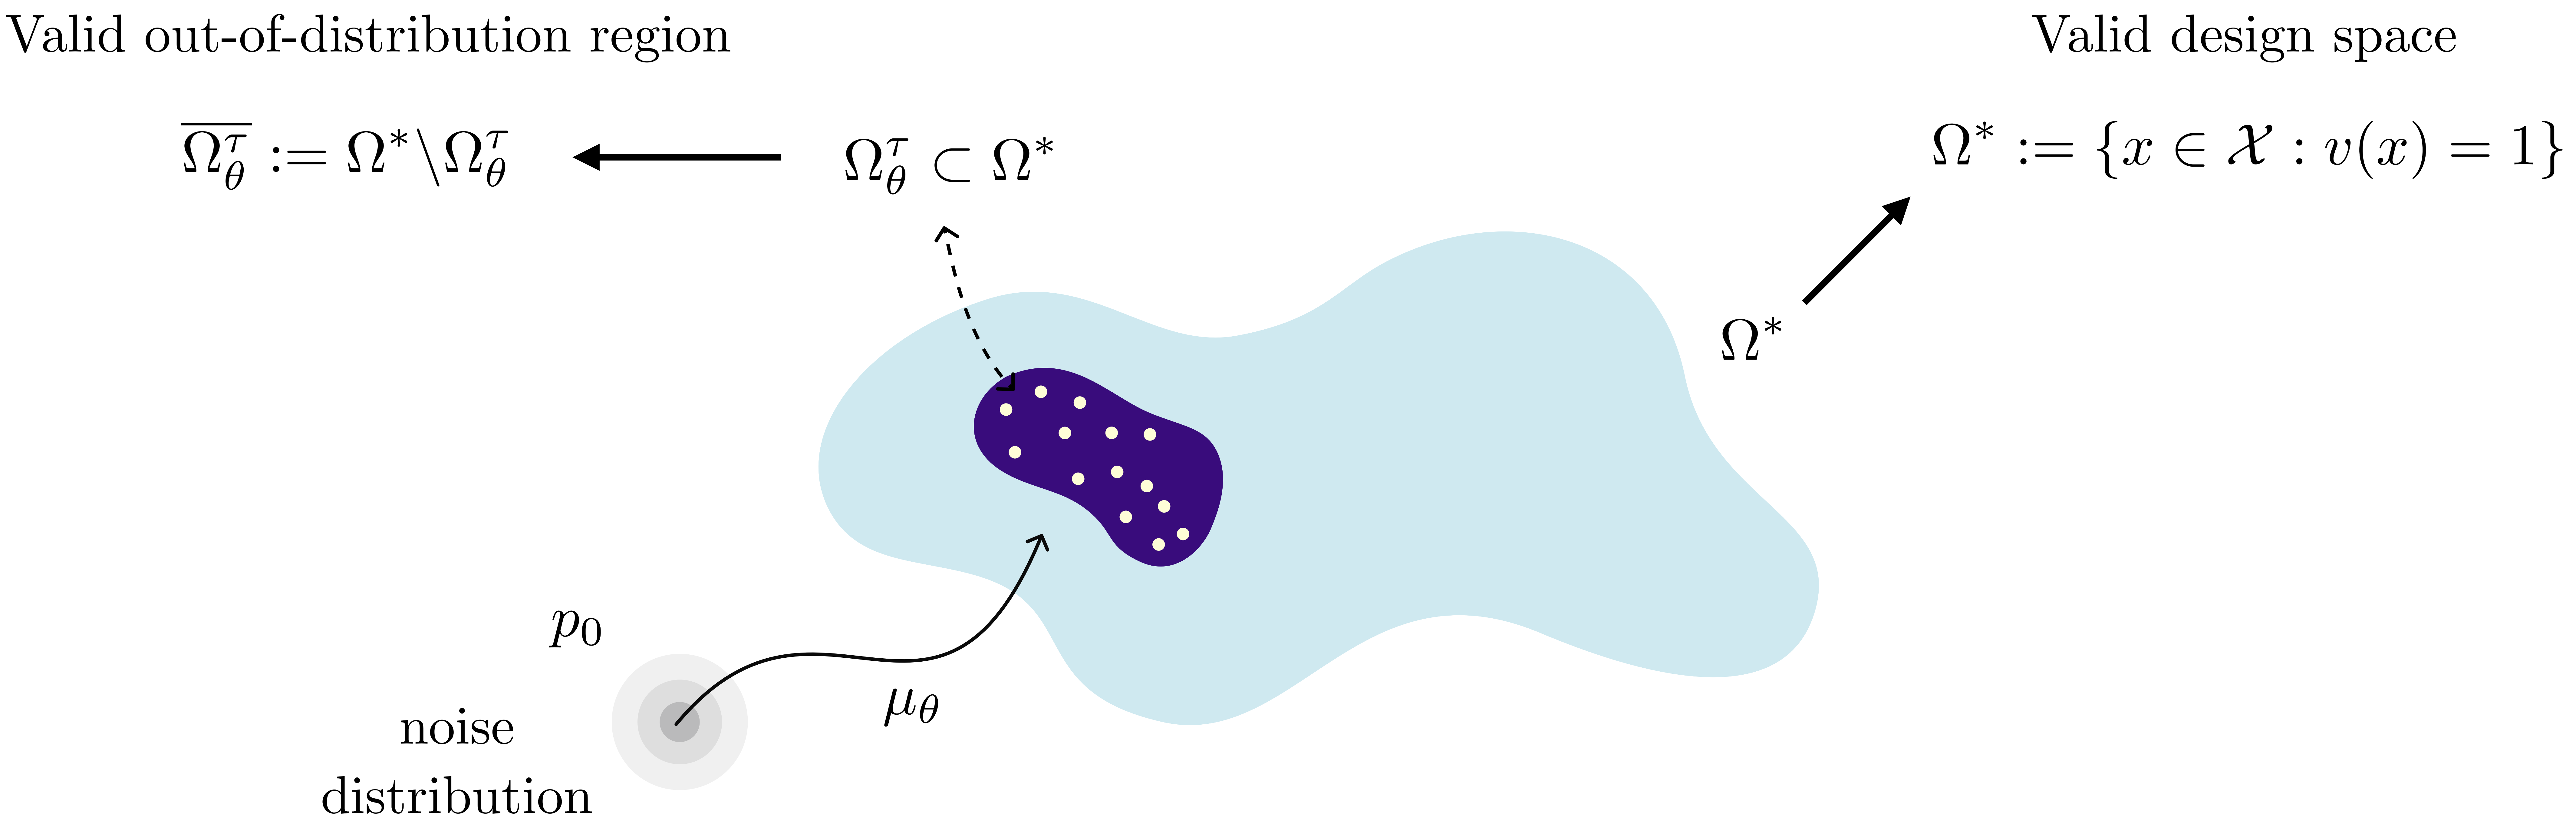

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

Generable Set vs Support

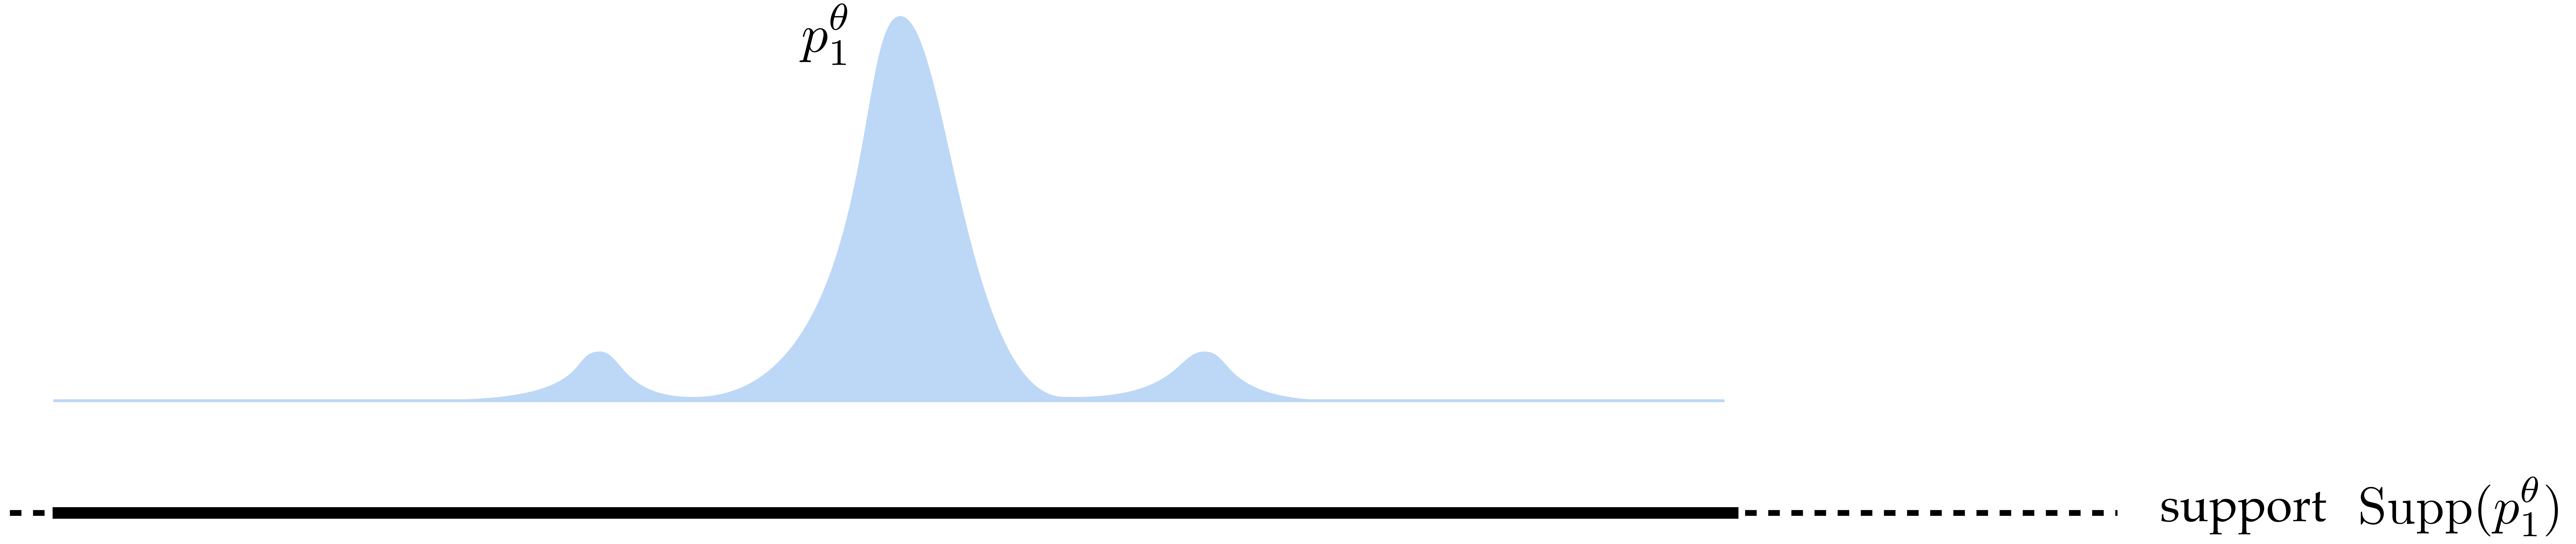

Generable Set vs Support

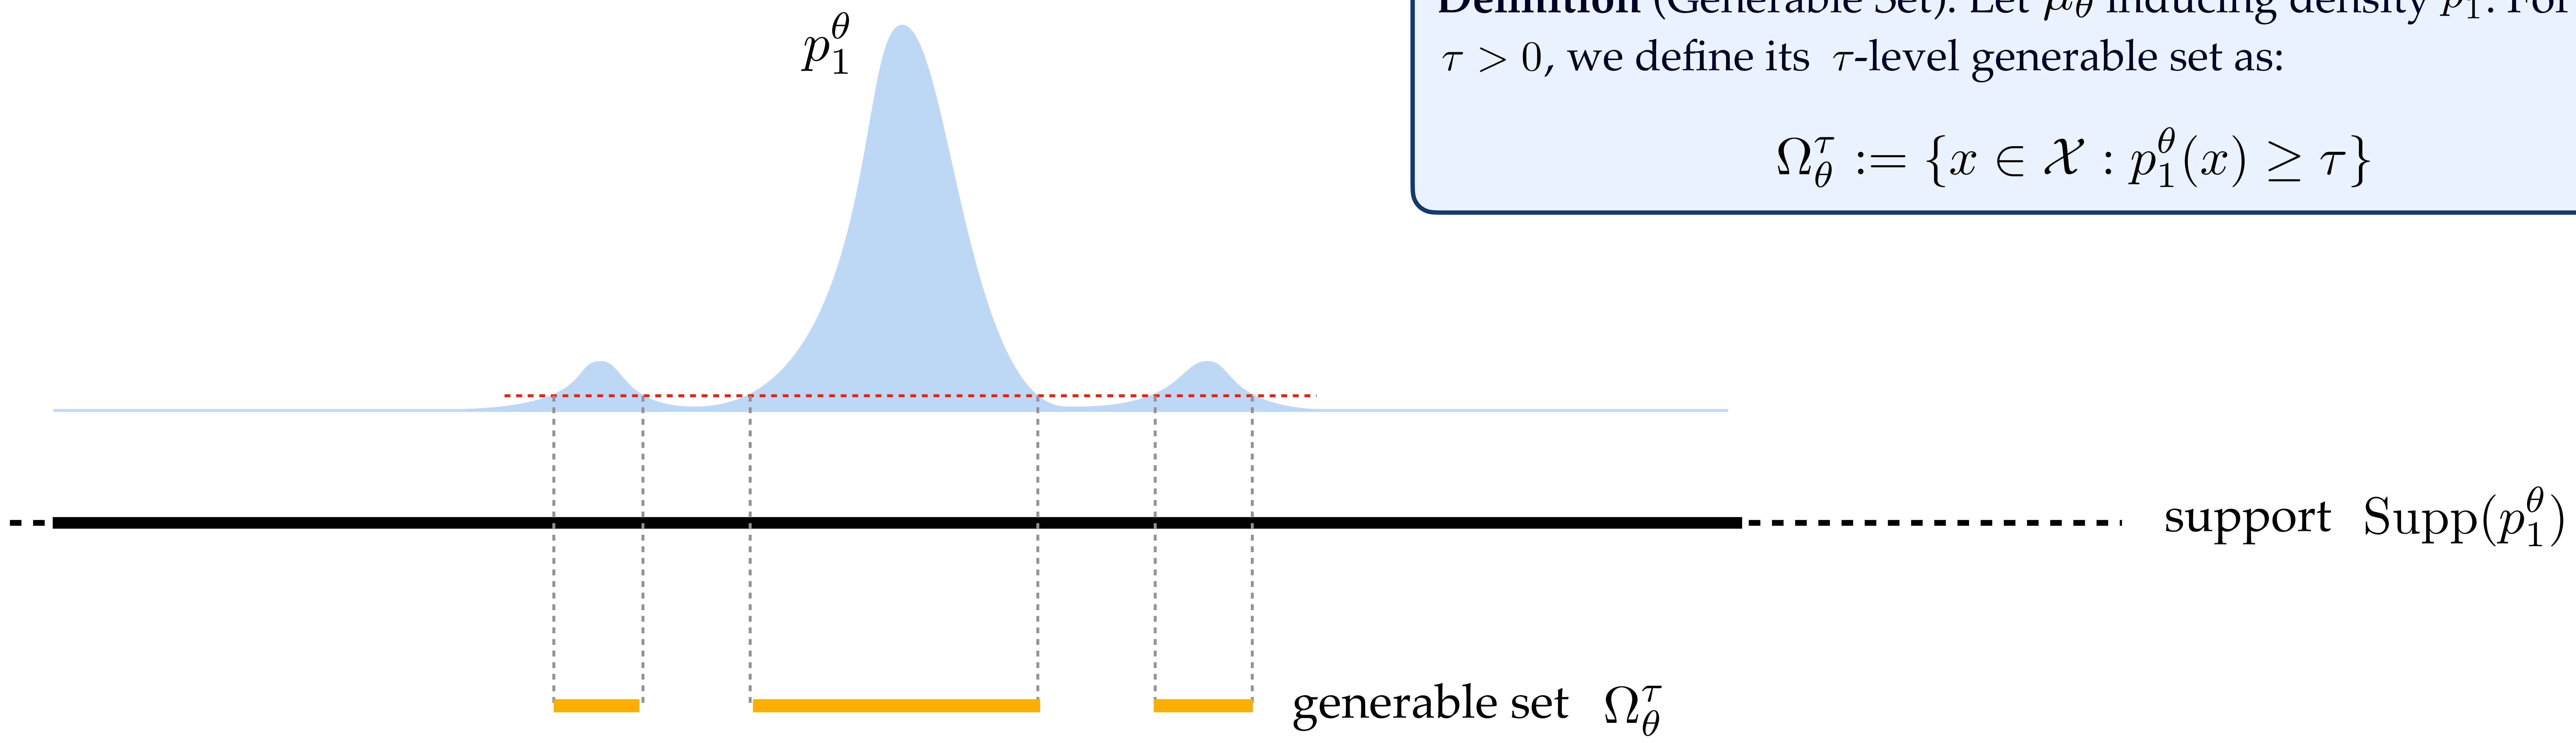

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

Generable Set vs Support

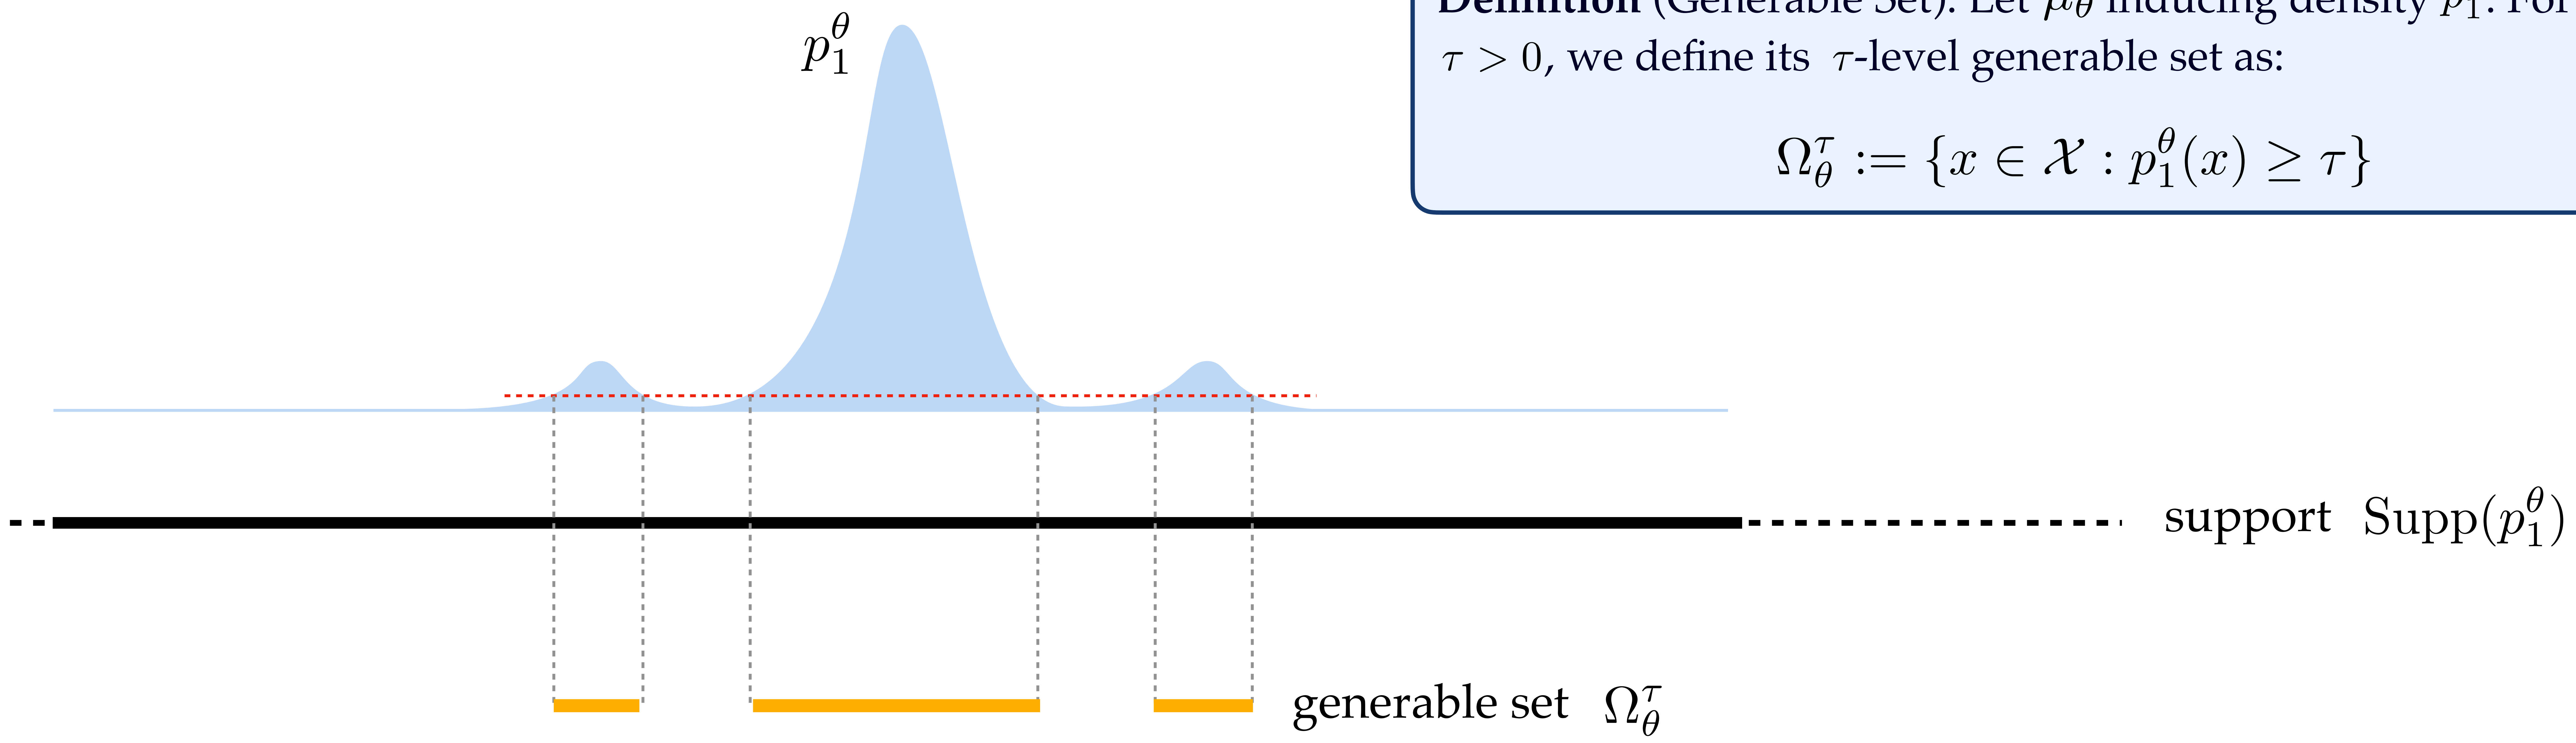

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

$$\Omega_\theta^\tau \xrightarrow{\tau \to 0} \text{Supp}(p_1^\theta)$$

Generable Set vs Support

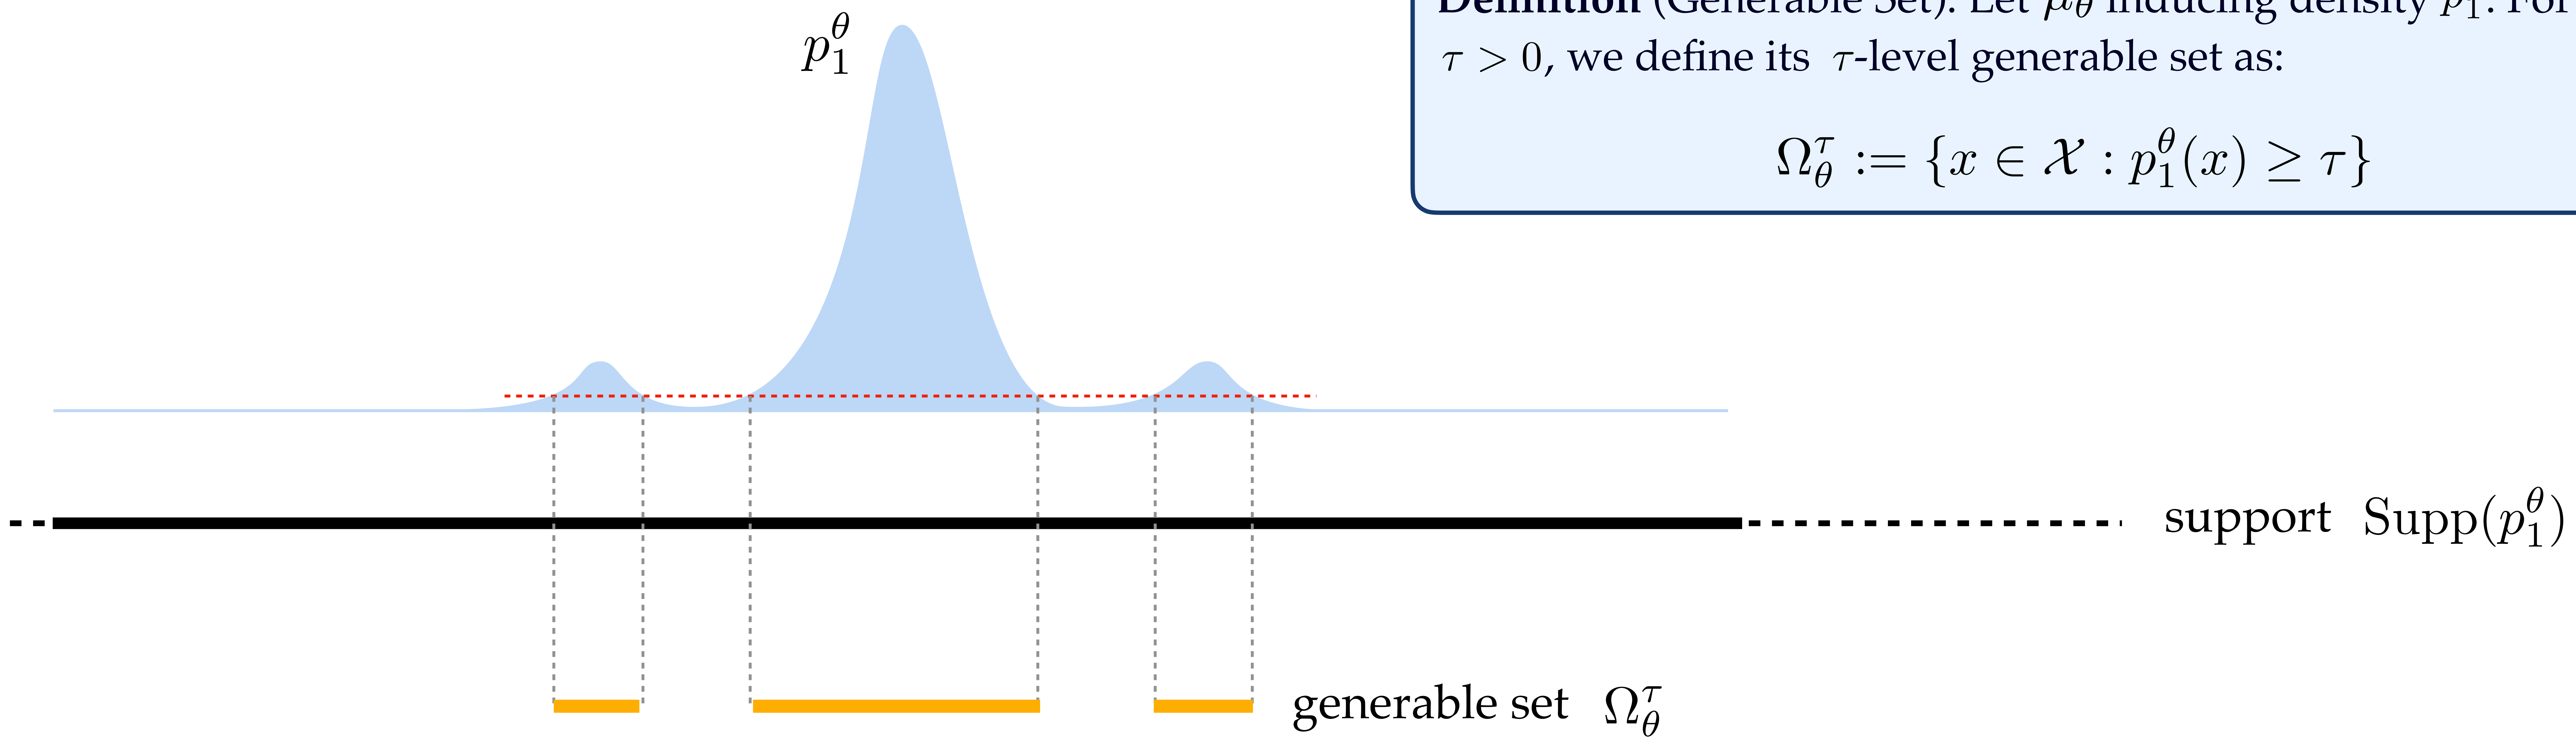

Definition (Generable Set). Let μ_θ inducing density p_1^θ. For $\tau > 0$, we define its τ-level generable set as:

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

generable set →
- region likely to be covered by a finite draw
- region where the generative model is reliable (e.g., sufficient validity signal)

Generable Set Expansion for OOD Flow Modeling

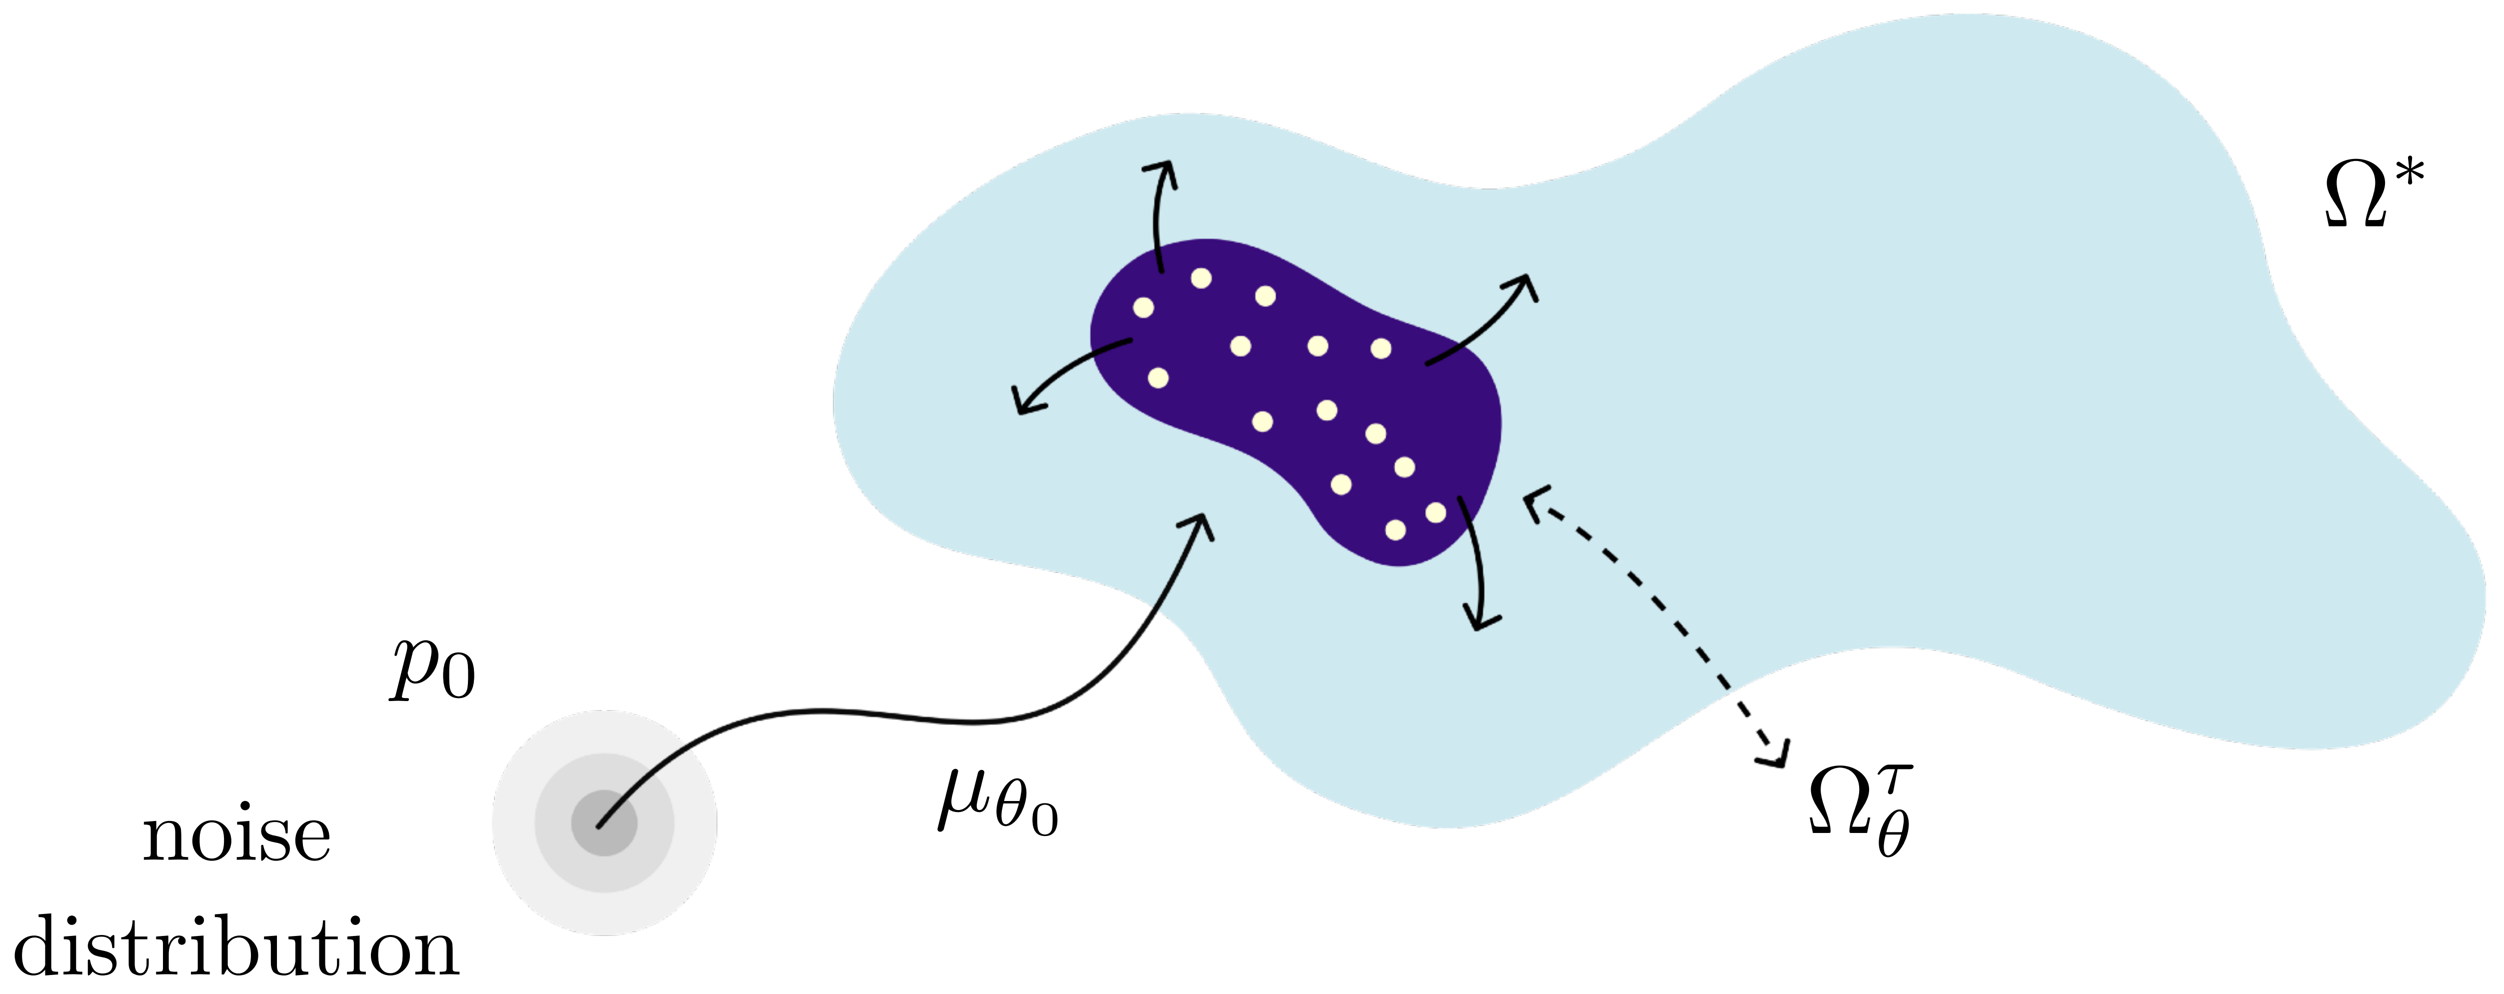

Generable Set Expansion: Out-of-Distribution Flow Modeling

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

Generable Set Expansion for OOD Flow Modeling

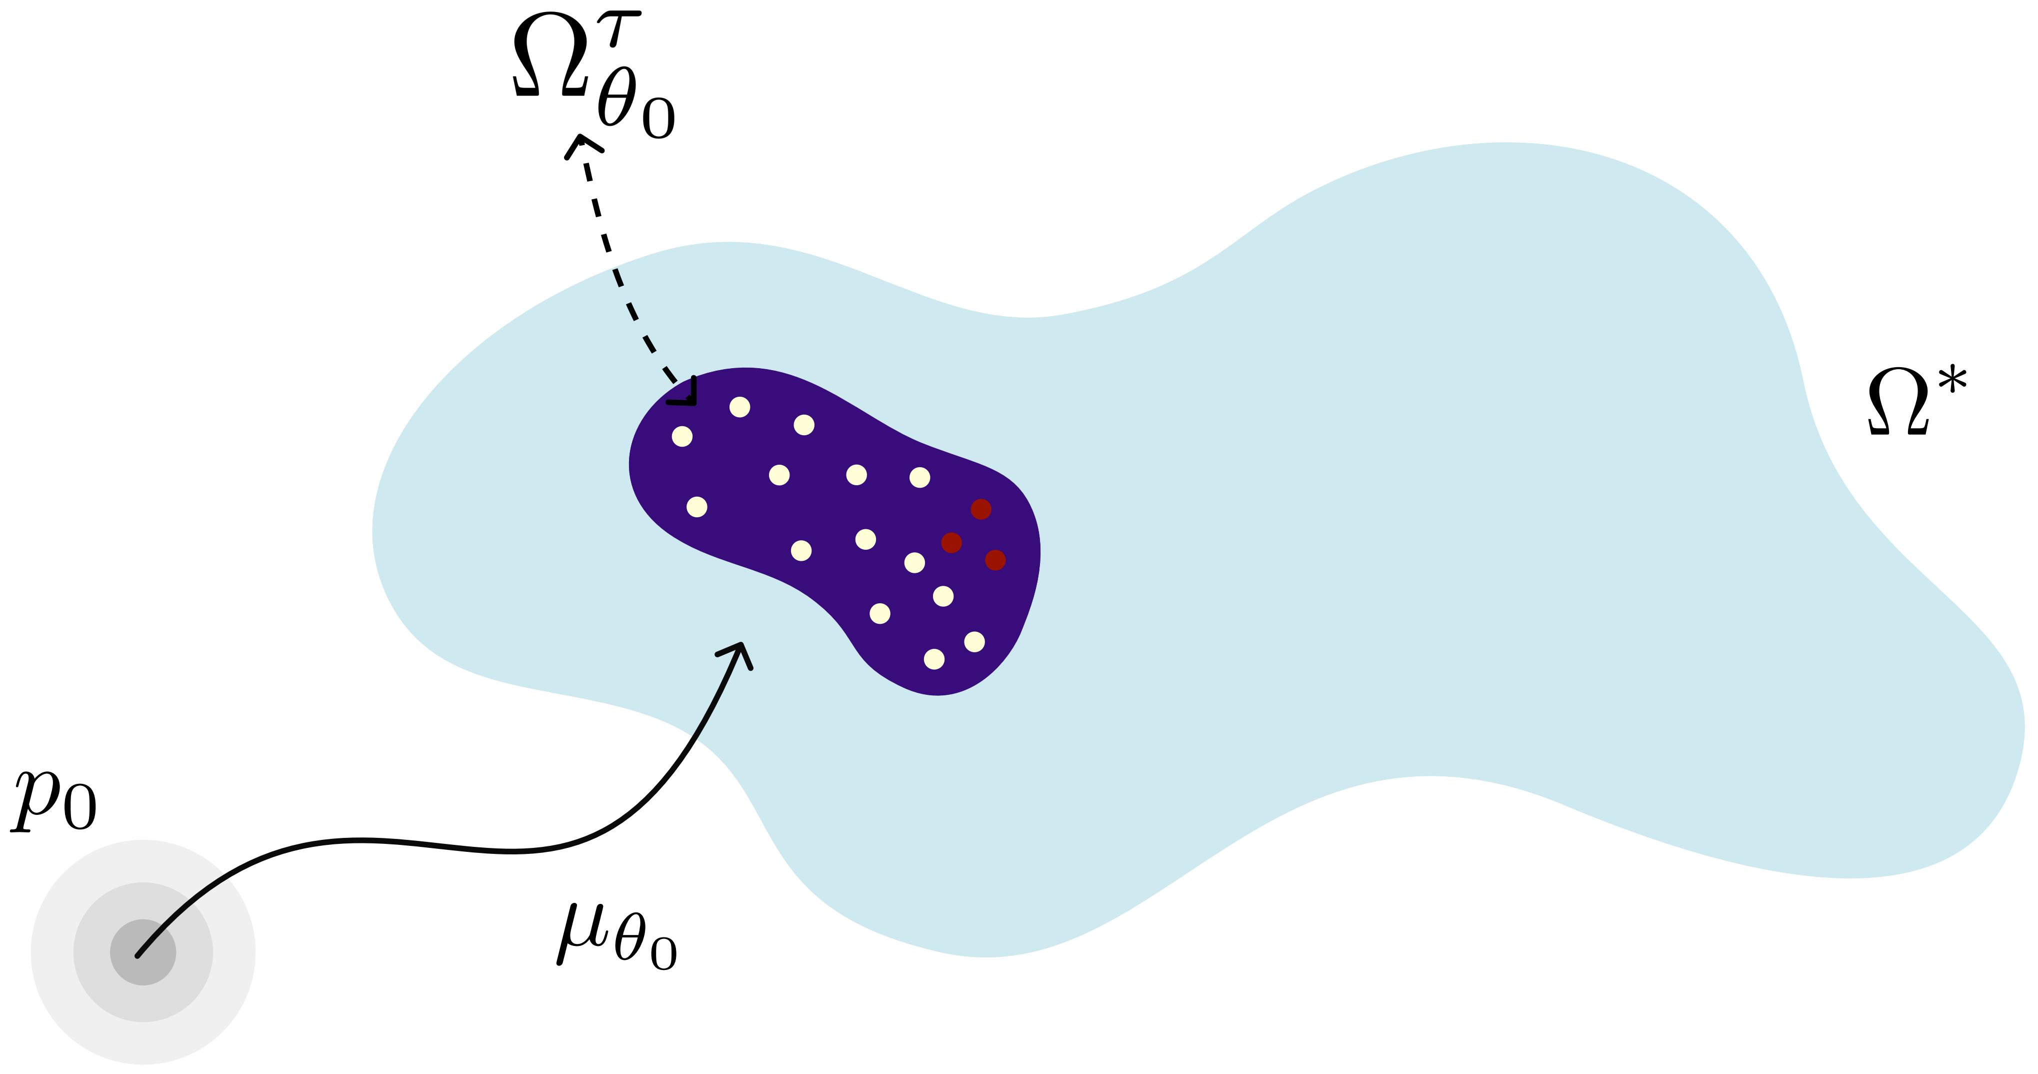

Black-box noisy binary verifier feedback

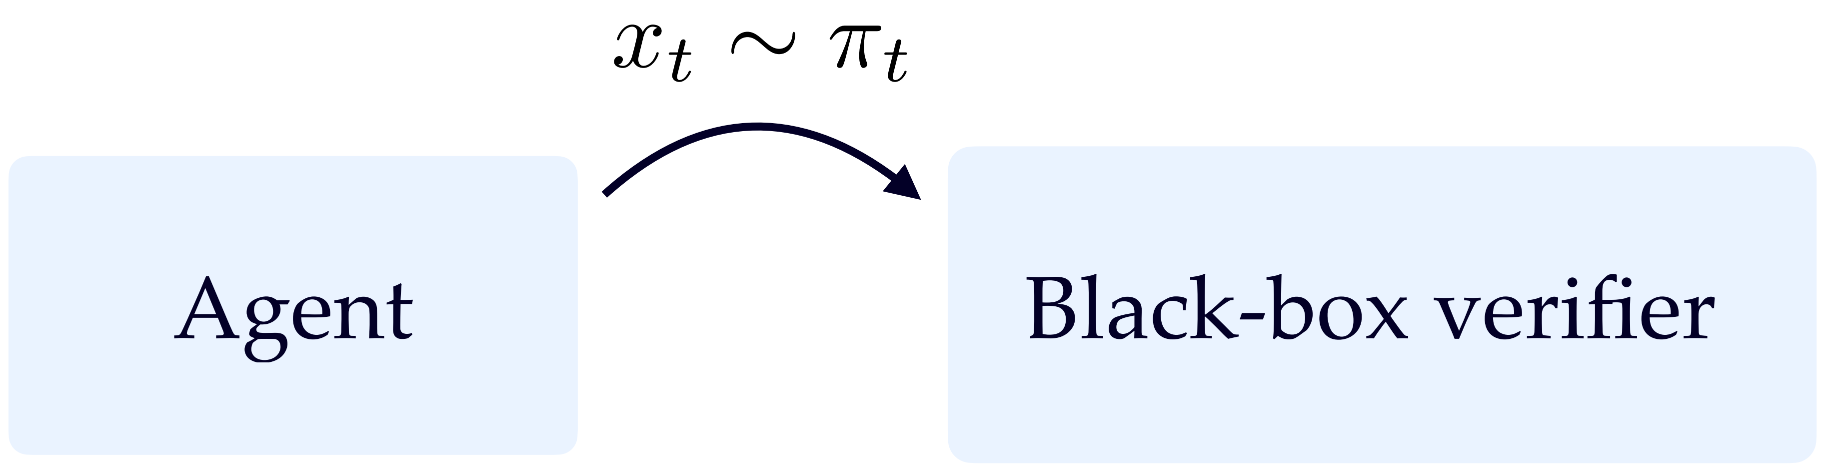

$$\theta^{\text{pre}} = \theta_0$$

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

Generable Set Expansion for OOD Flow Modeling

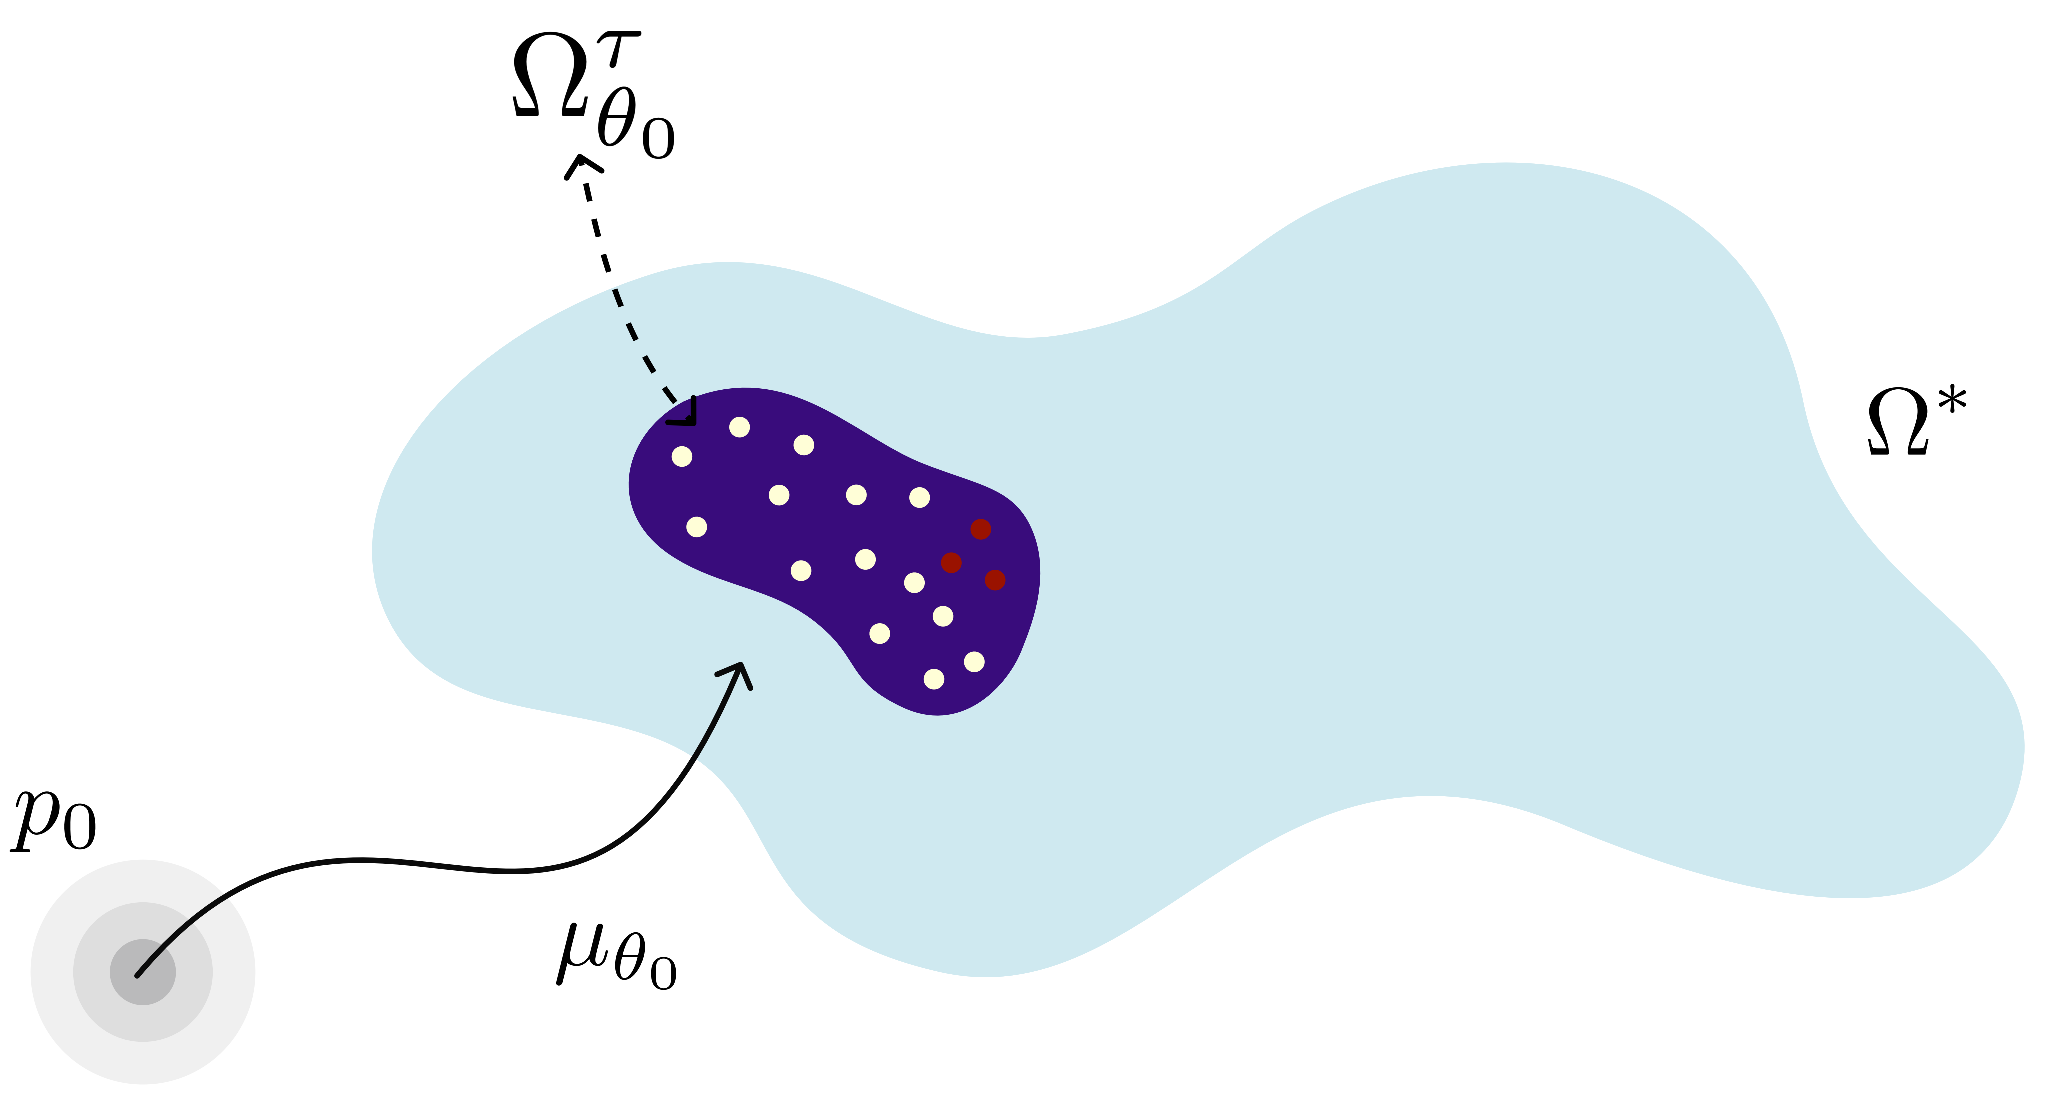

Black-box noisy binary verifier feedback

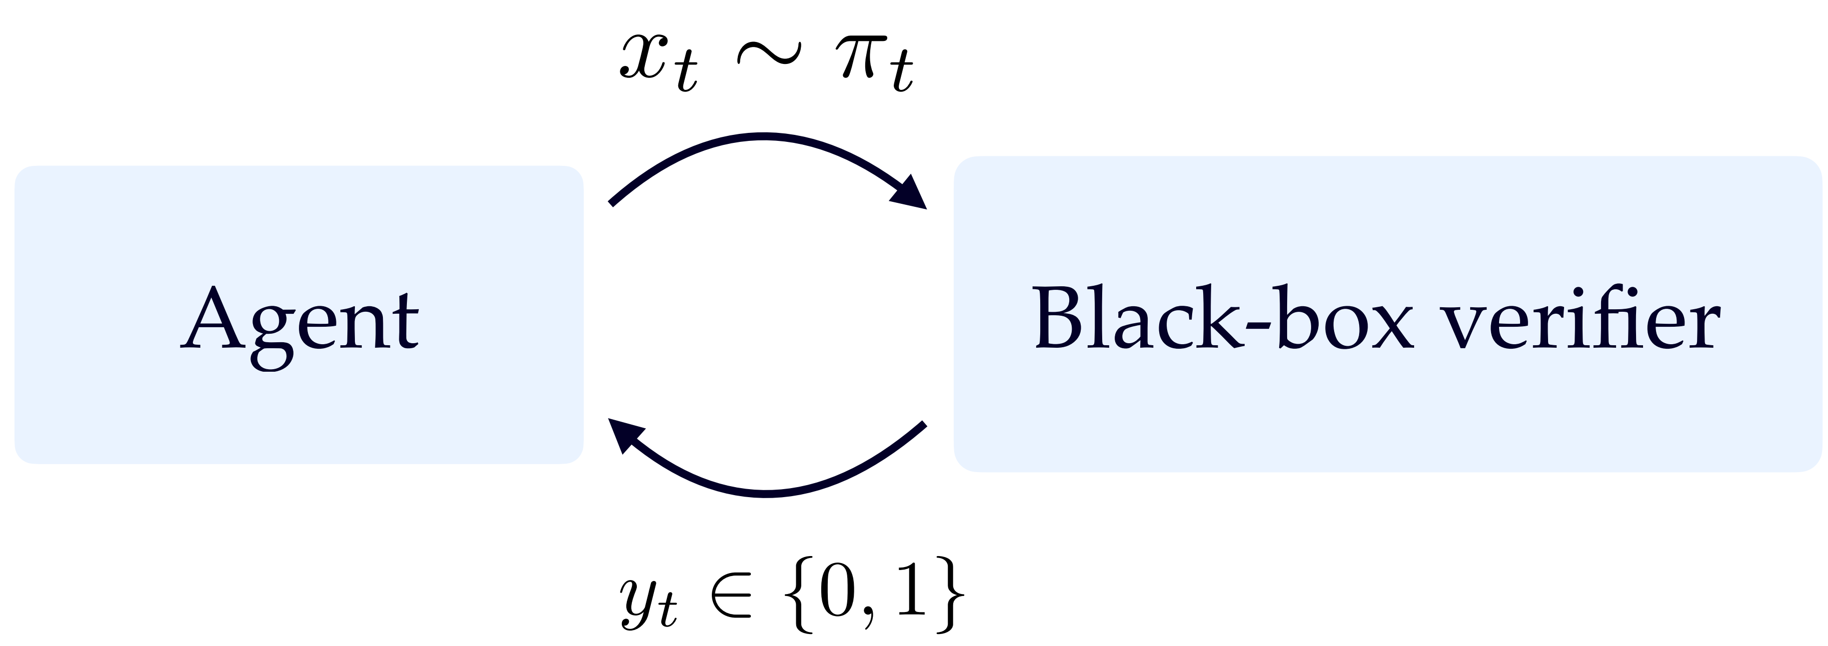

$\theta^{\text{pre}} = \theta_0$

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

Generable Set Expansion for OOD Flow Modeling

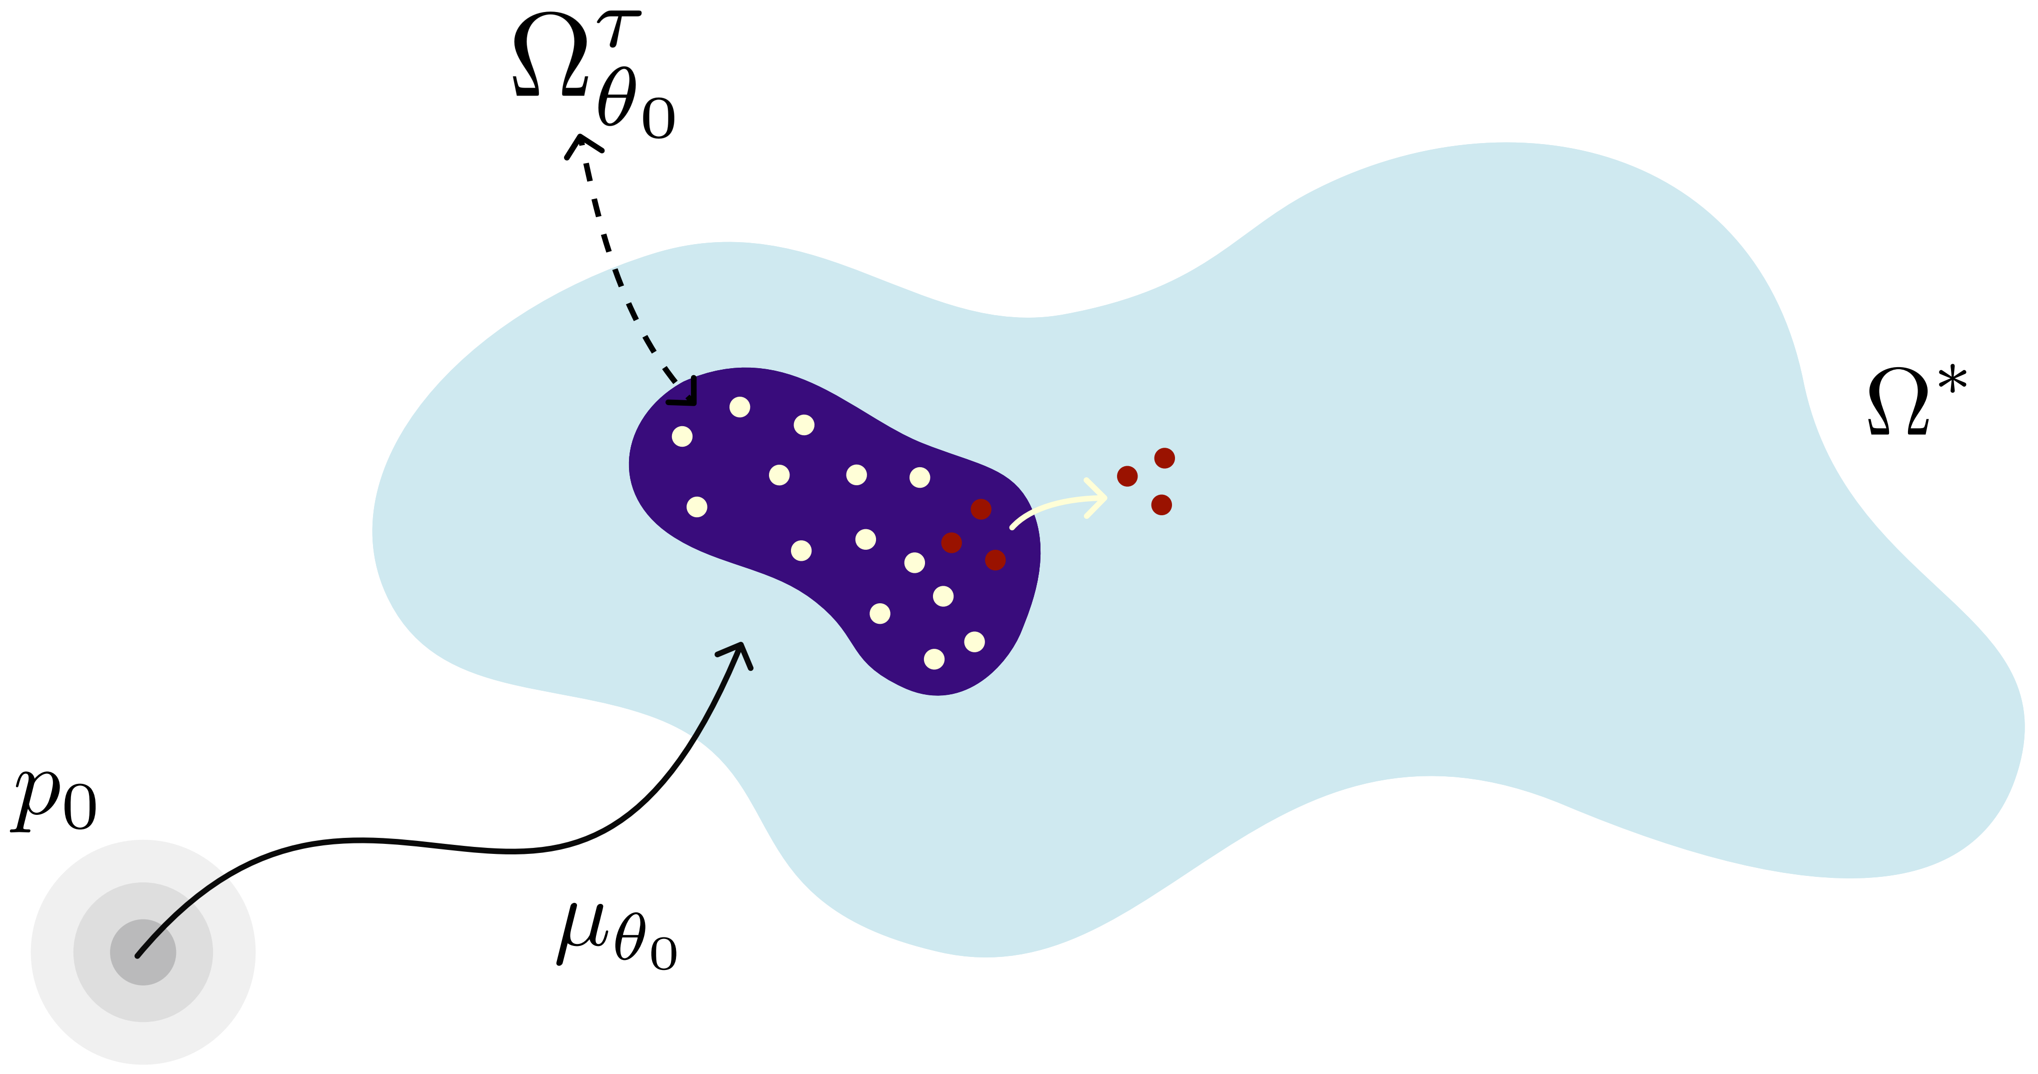

Black-box noisy binary verifier feedback

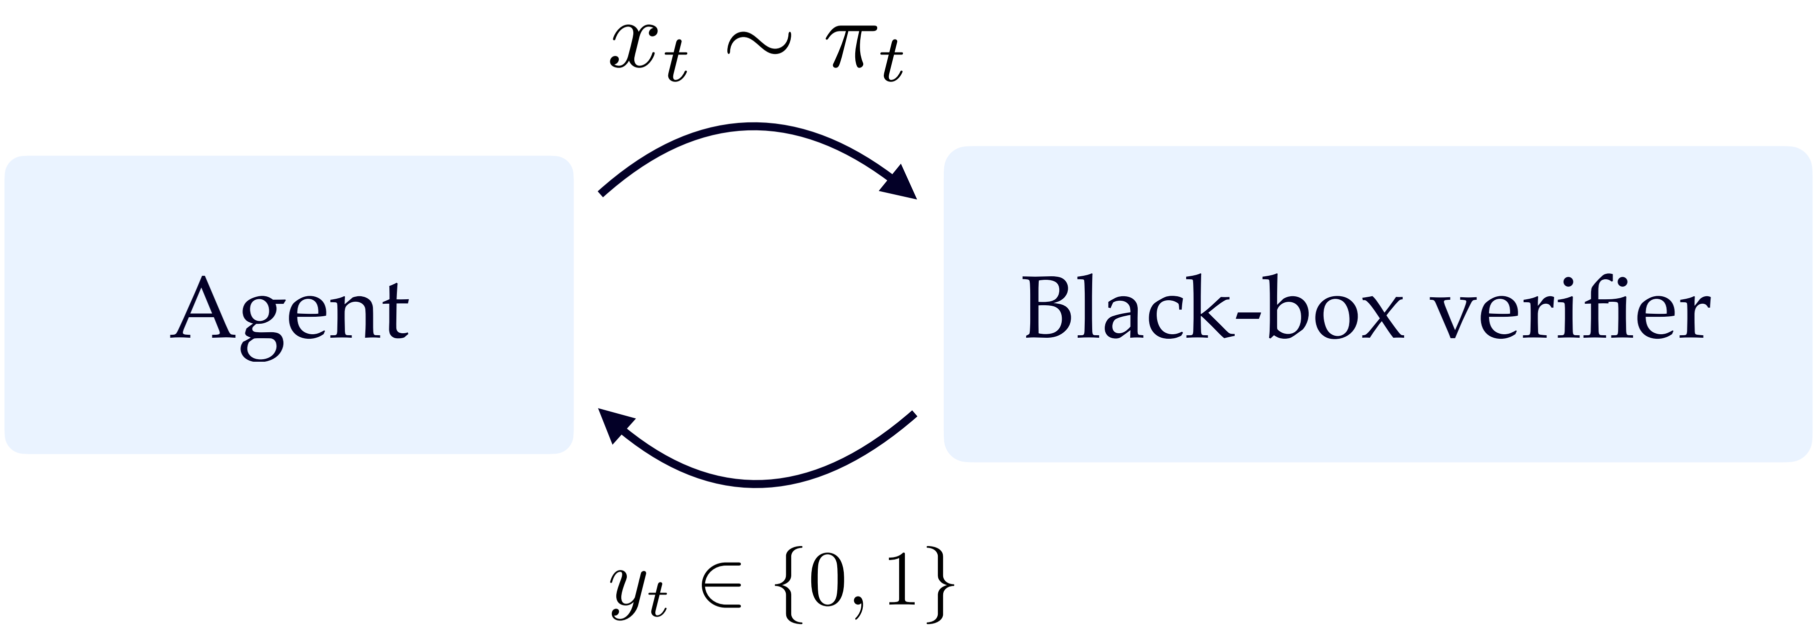

$\theta^{\text{pre}} = \theta_0$

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

Generable Set Expansion for OOD Flow Modeling

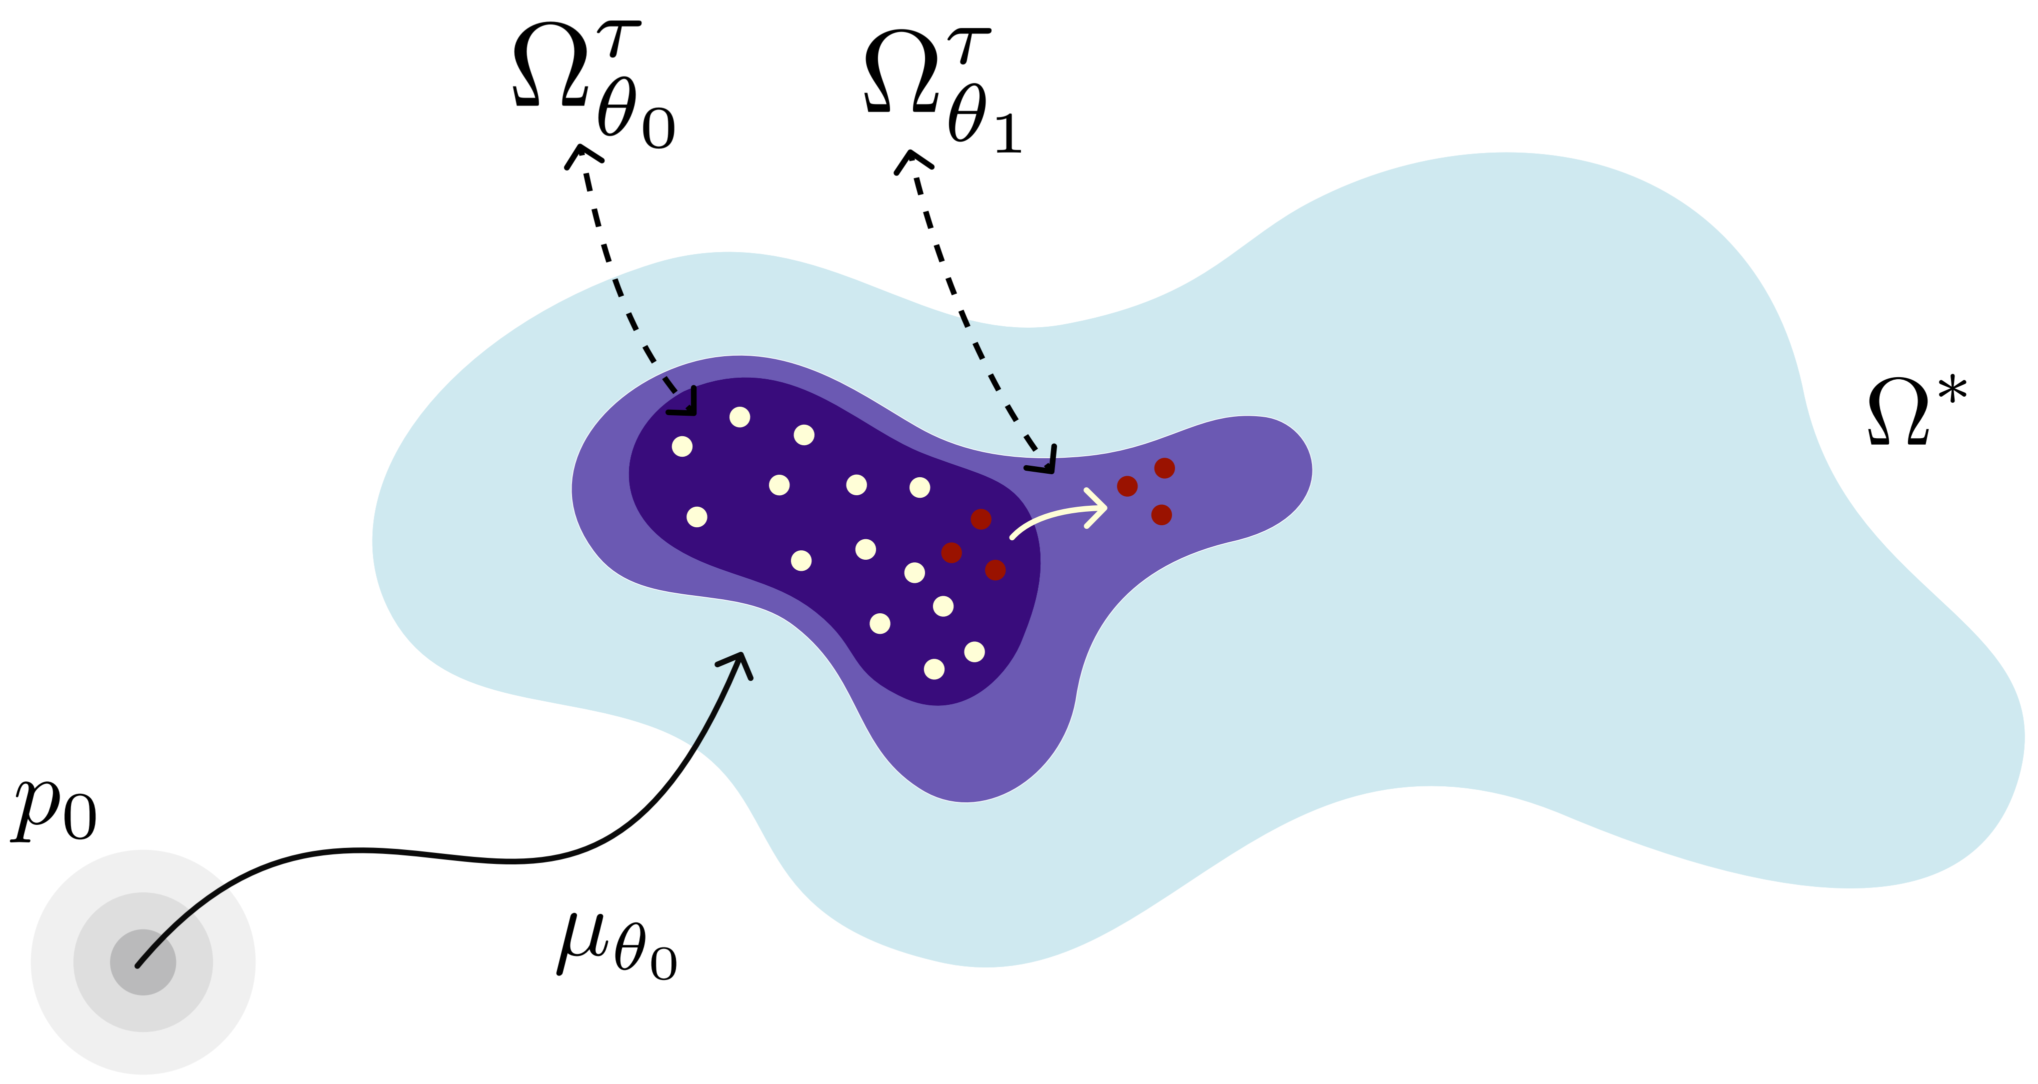

Black-box noisy binary verifier feedback

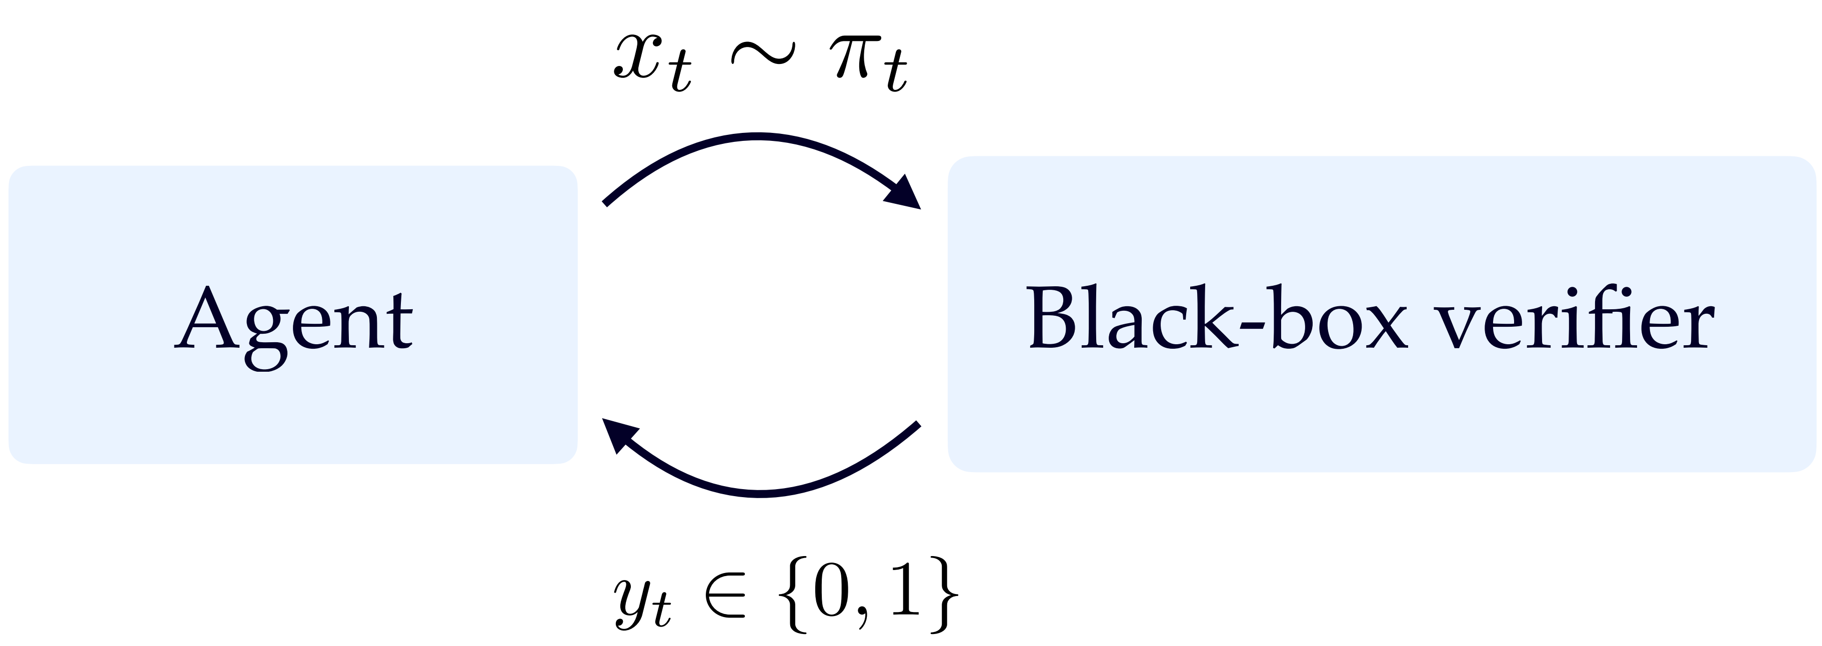

$$\theta^{\text{pre}} = \theta_0 \to \theta_1$$

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

Generable Set Expansion for OOD Flow Modeling

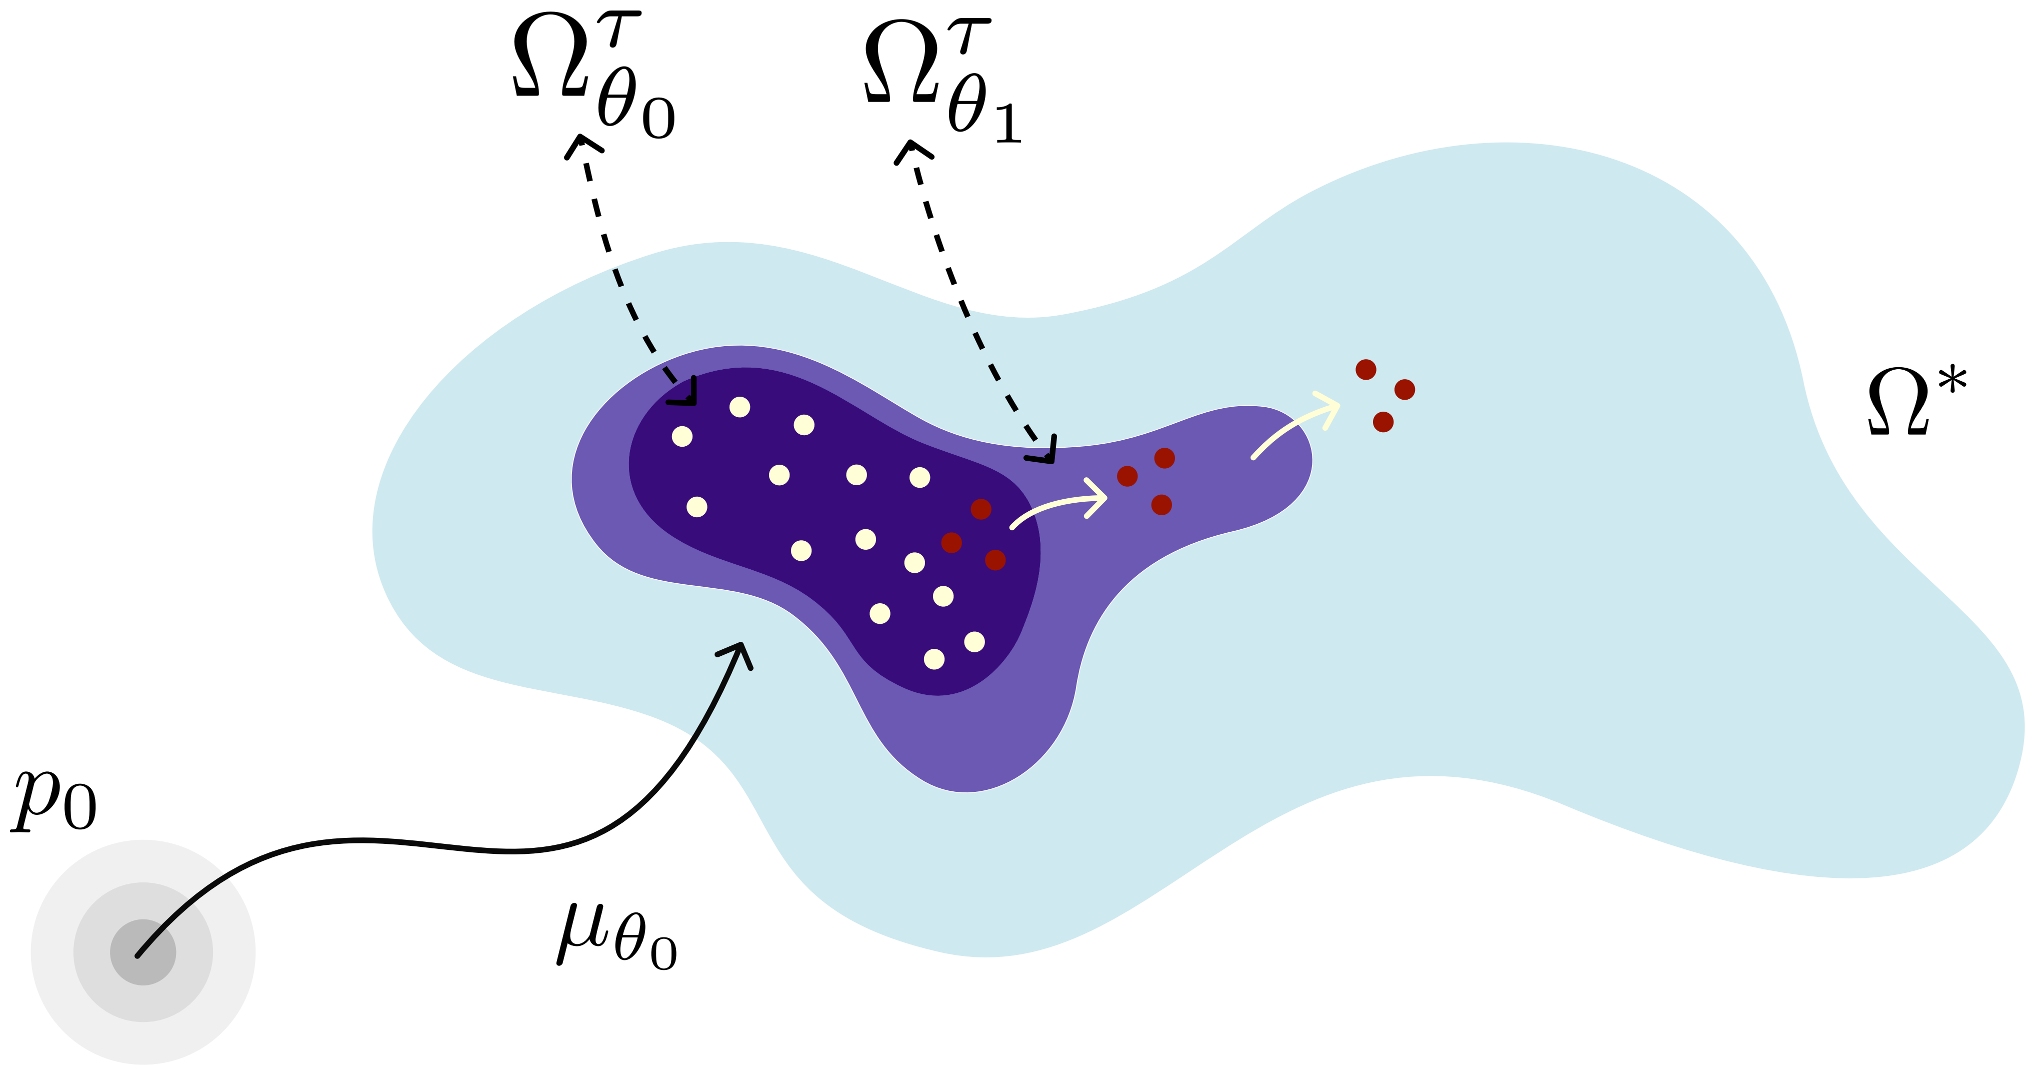

Black-box noisy binary verifier feedback

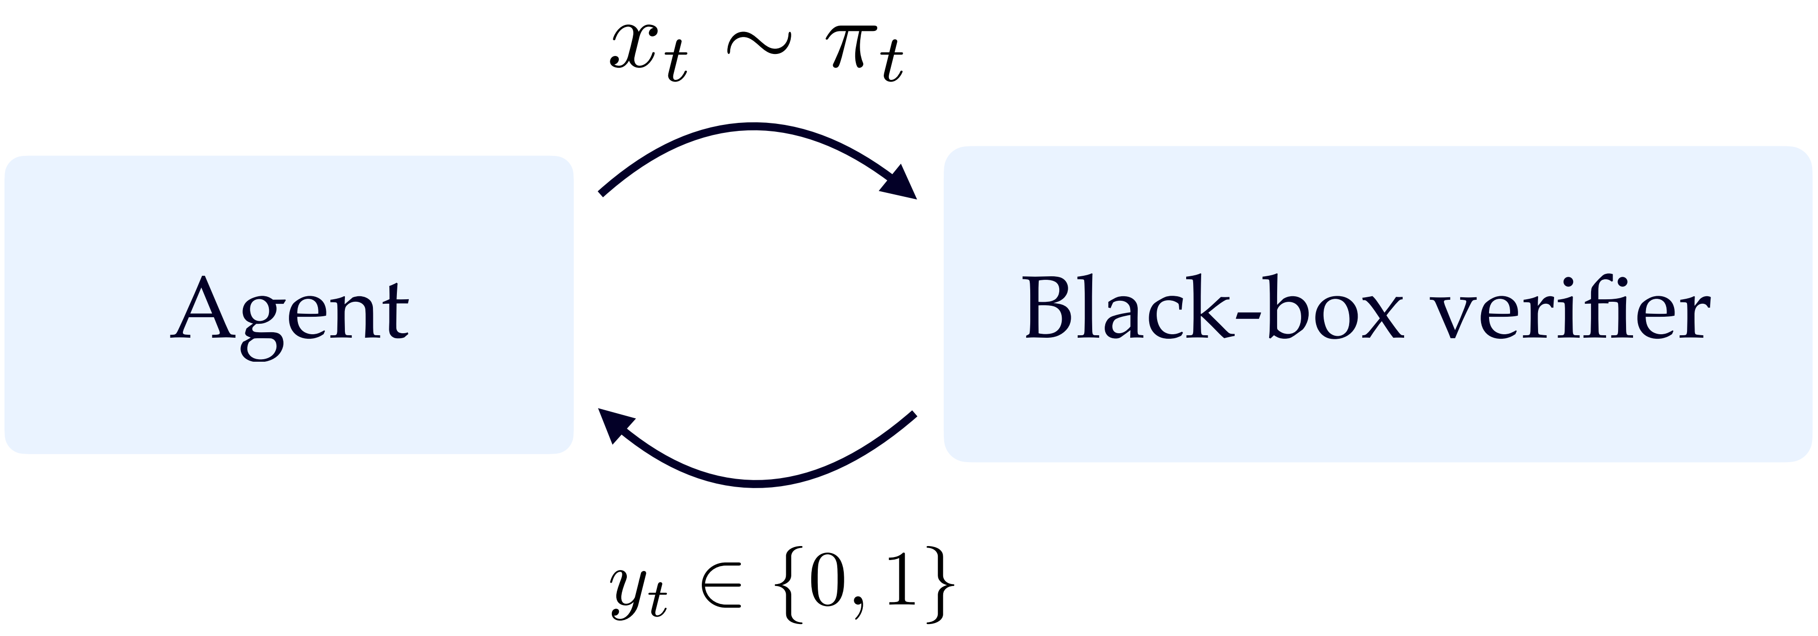

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to$$

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

Generable Set Expansion for OOD Flow Modeling

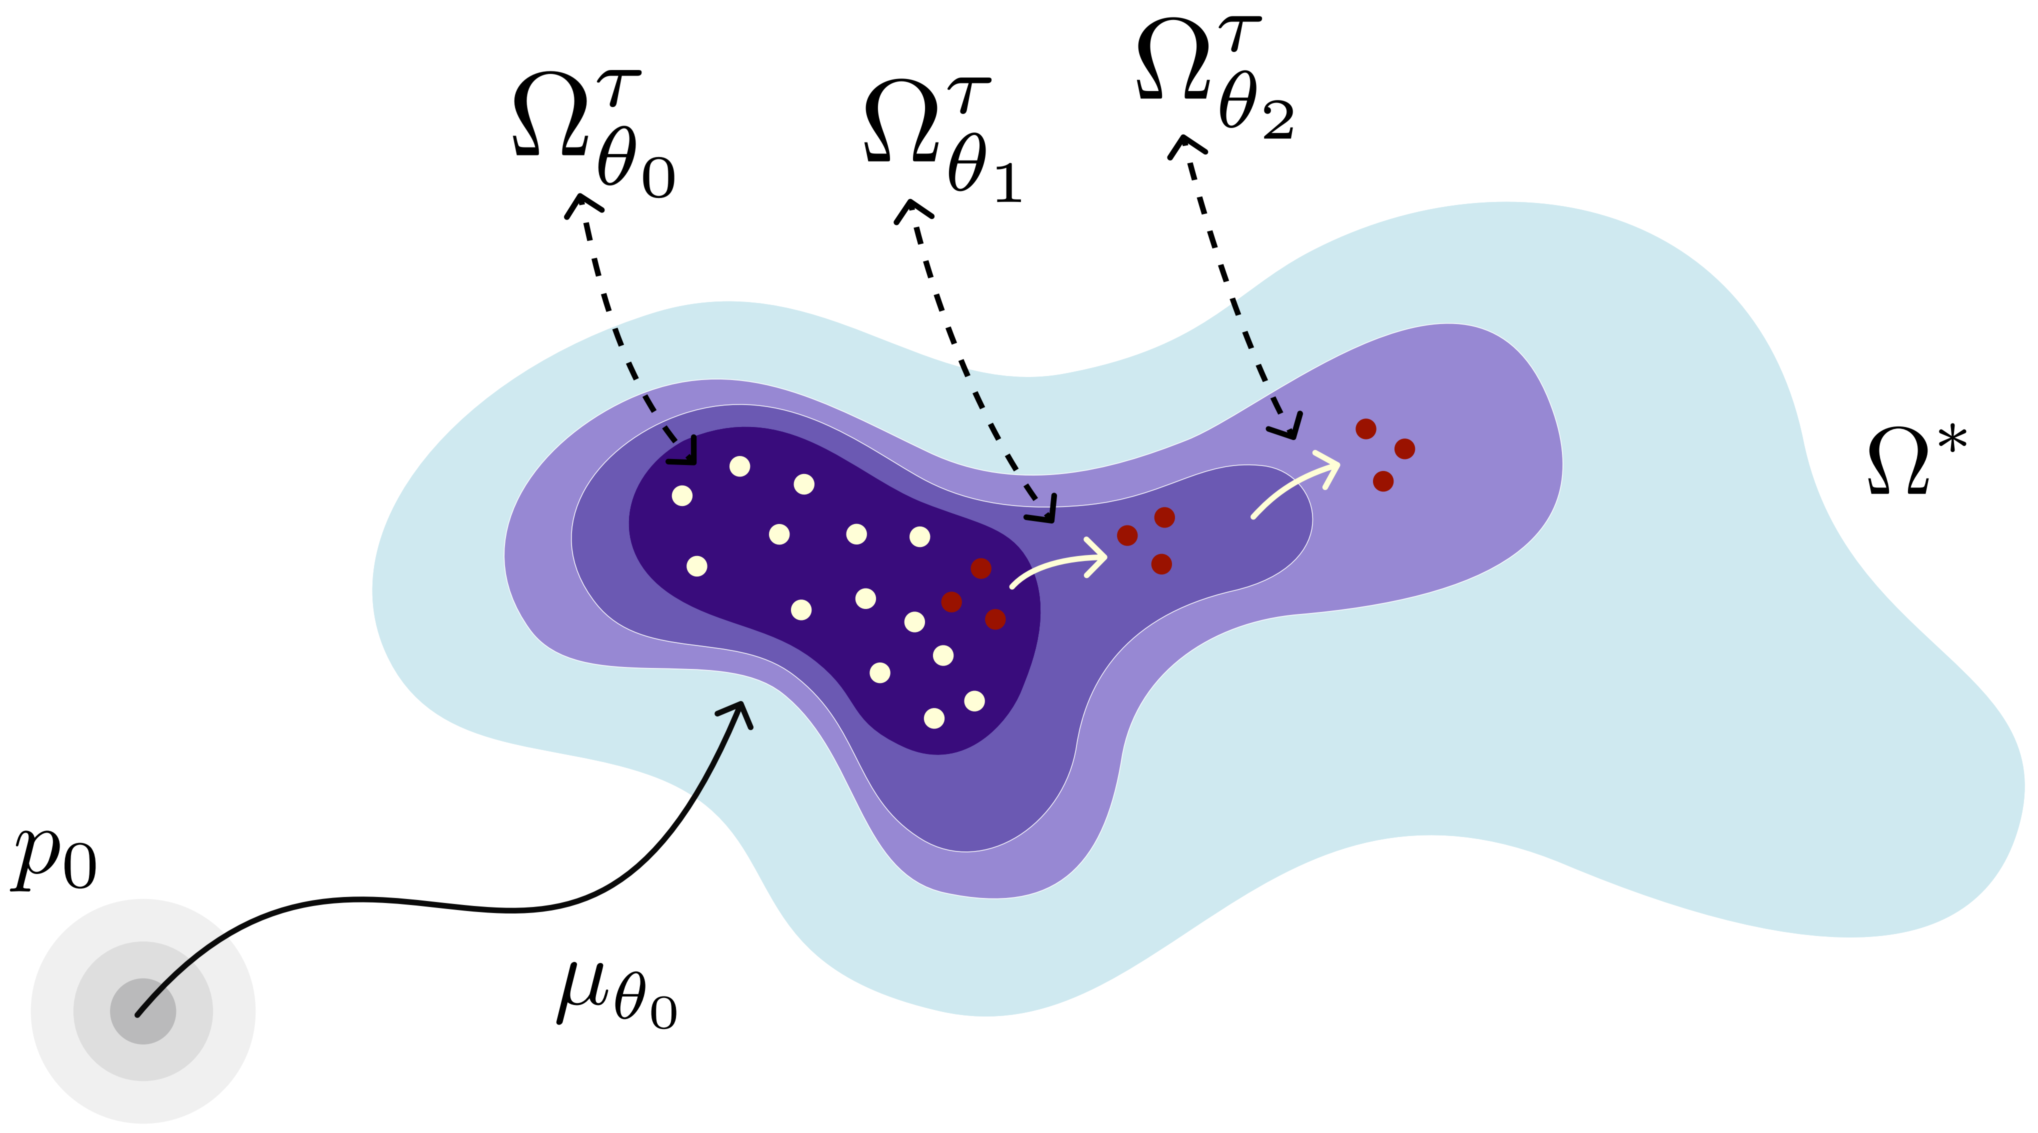

Black-box noisy binary verifier feedback

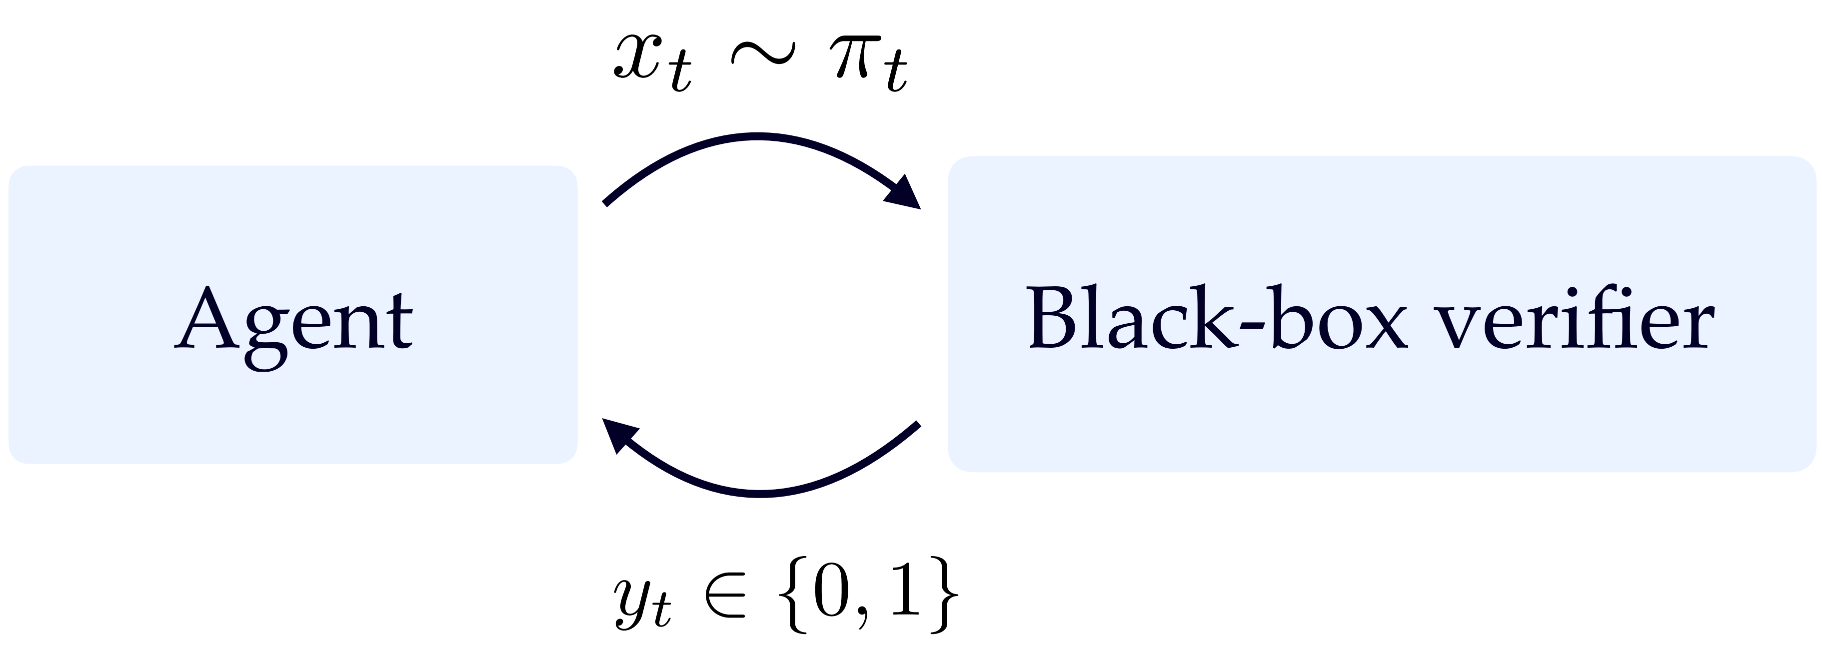

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots$$

Generable Set Expansion for OOD Flow Modeling

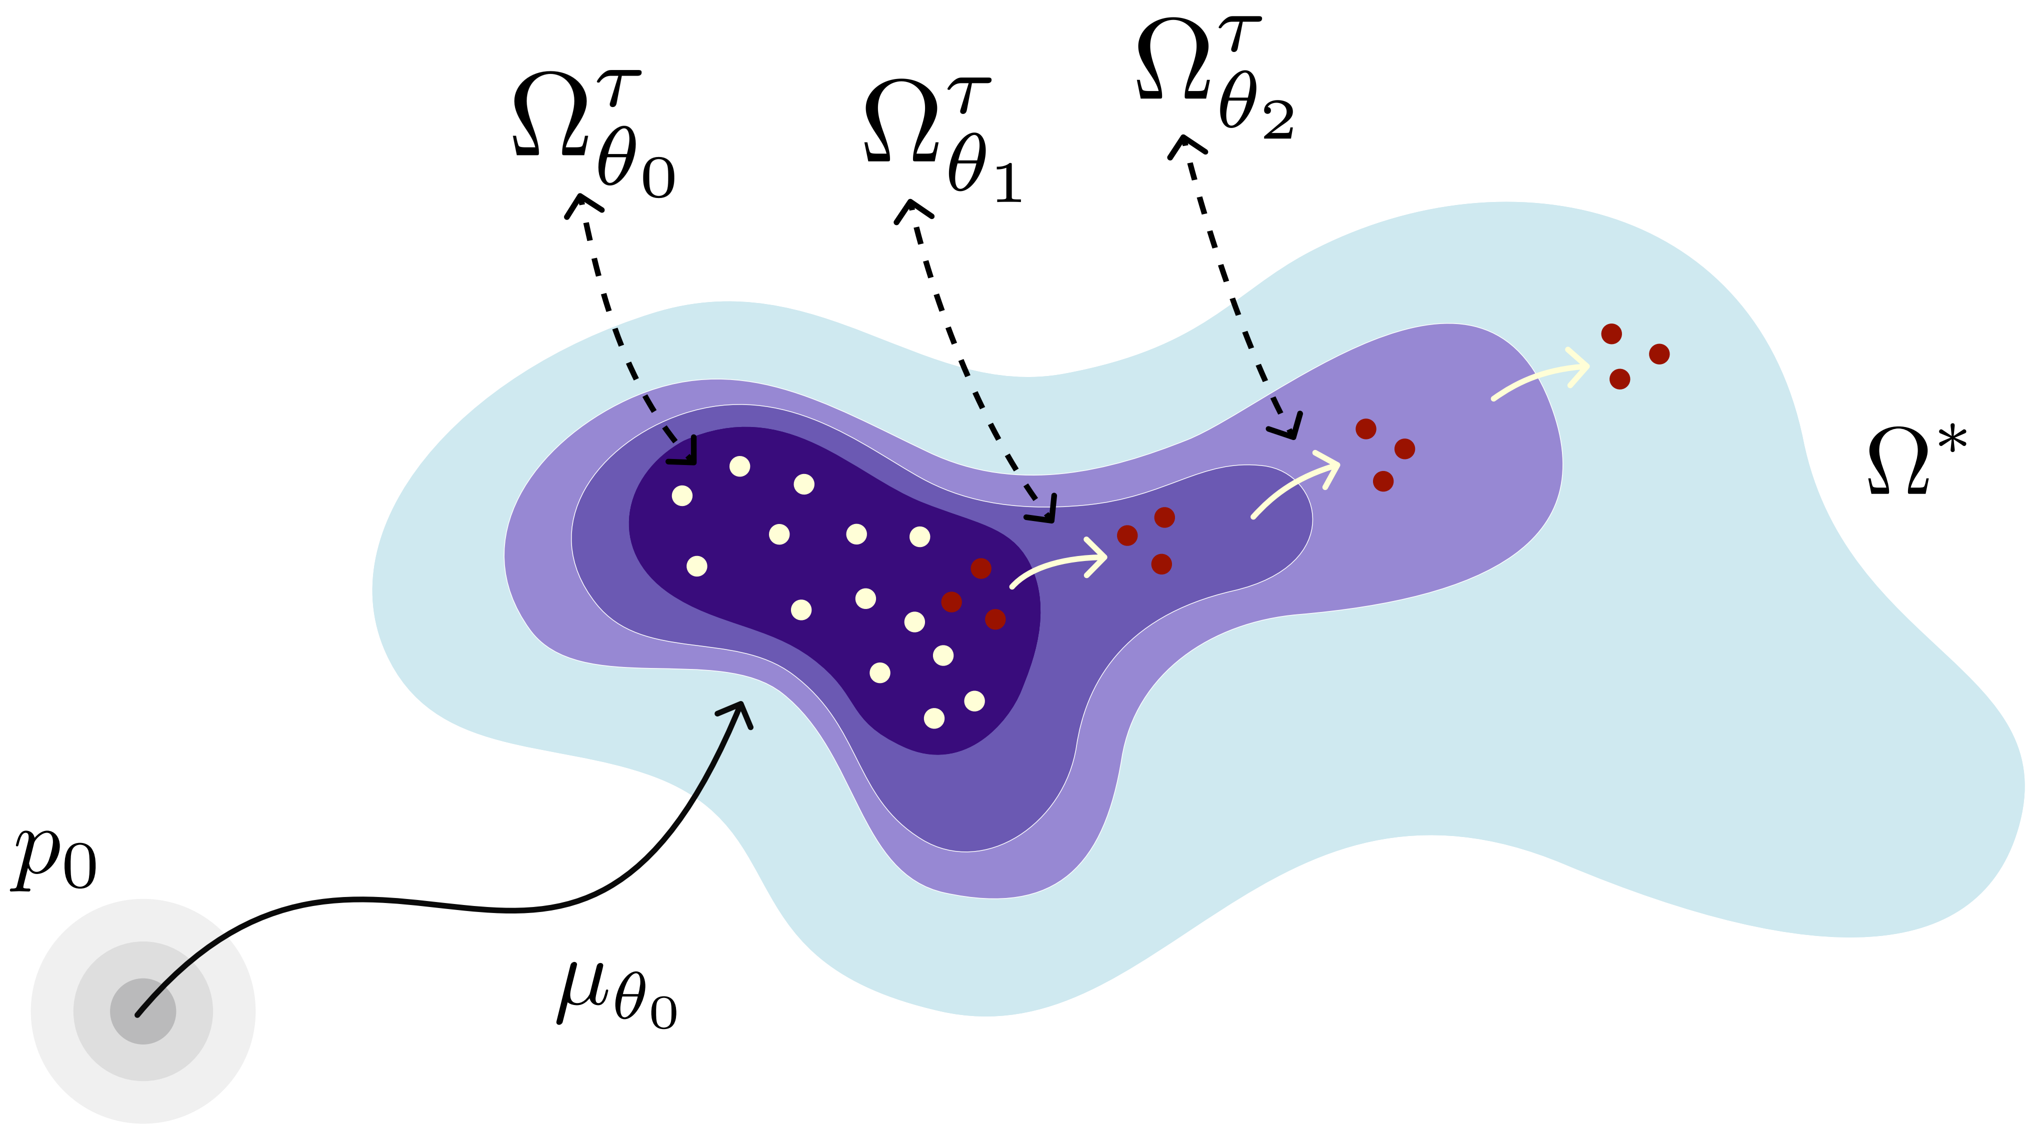

Black-box noisy binary verifier feedback

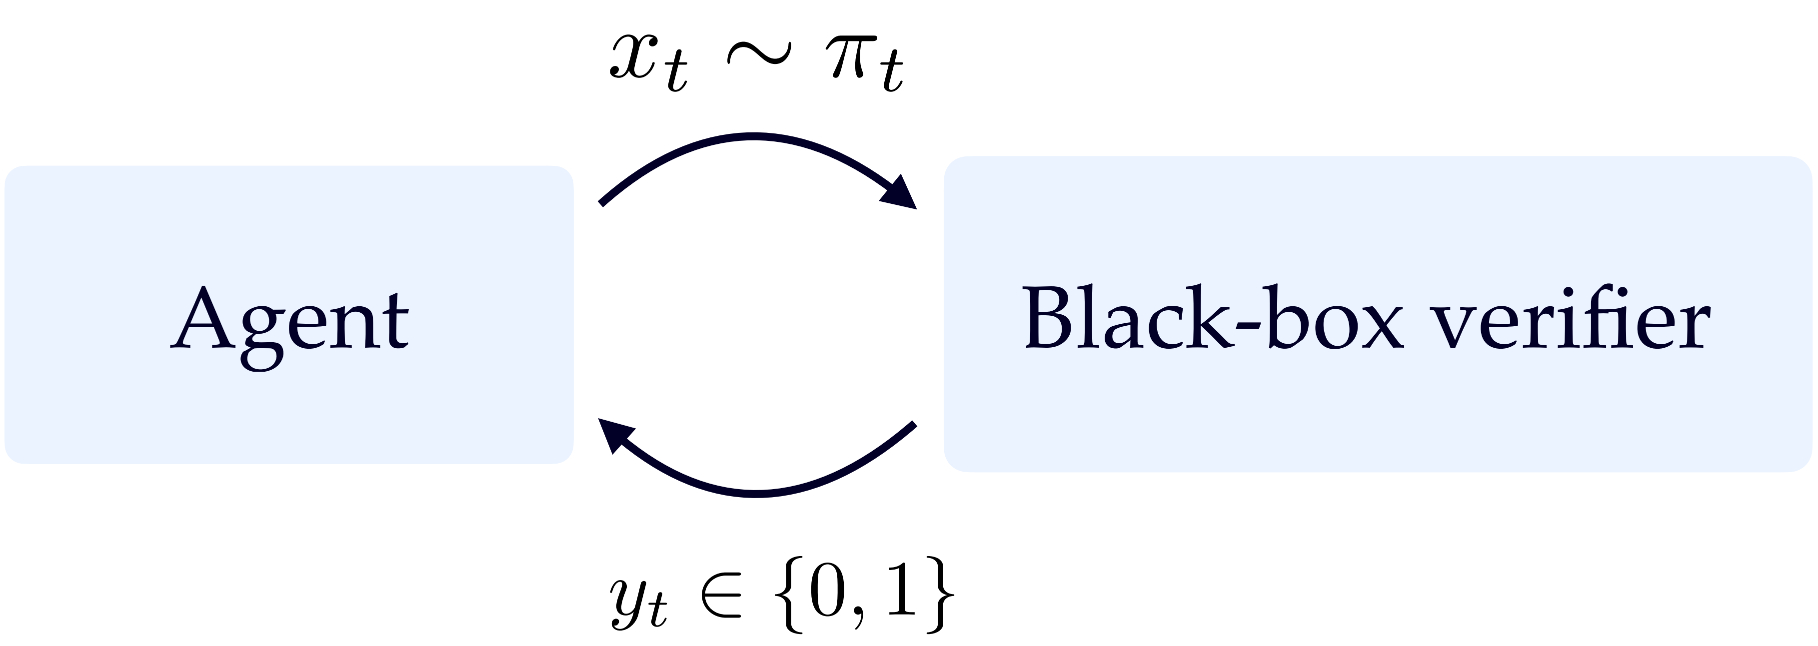

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to$$

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

Generable Set Expansion for OOD Flow Modeling

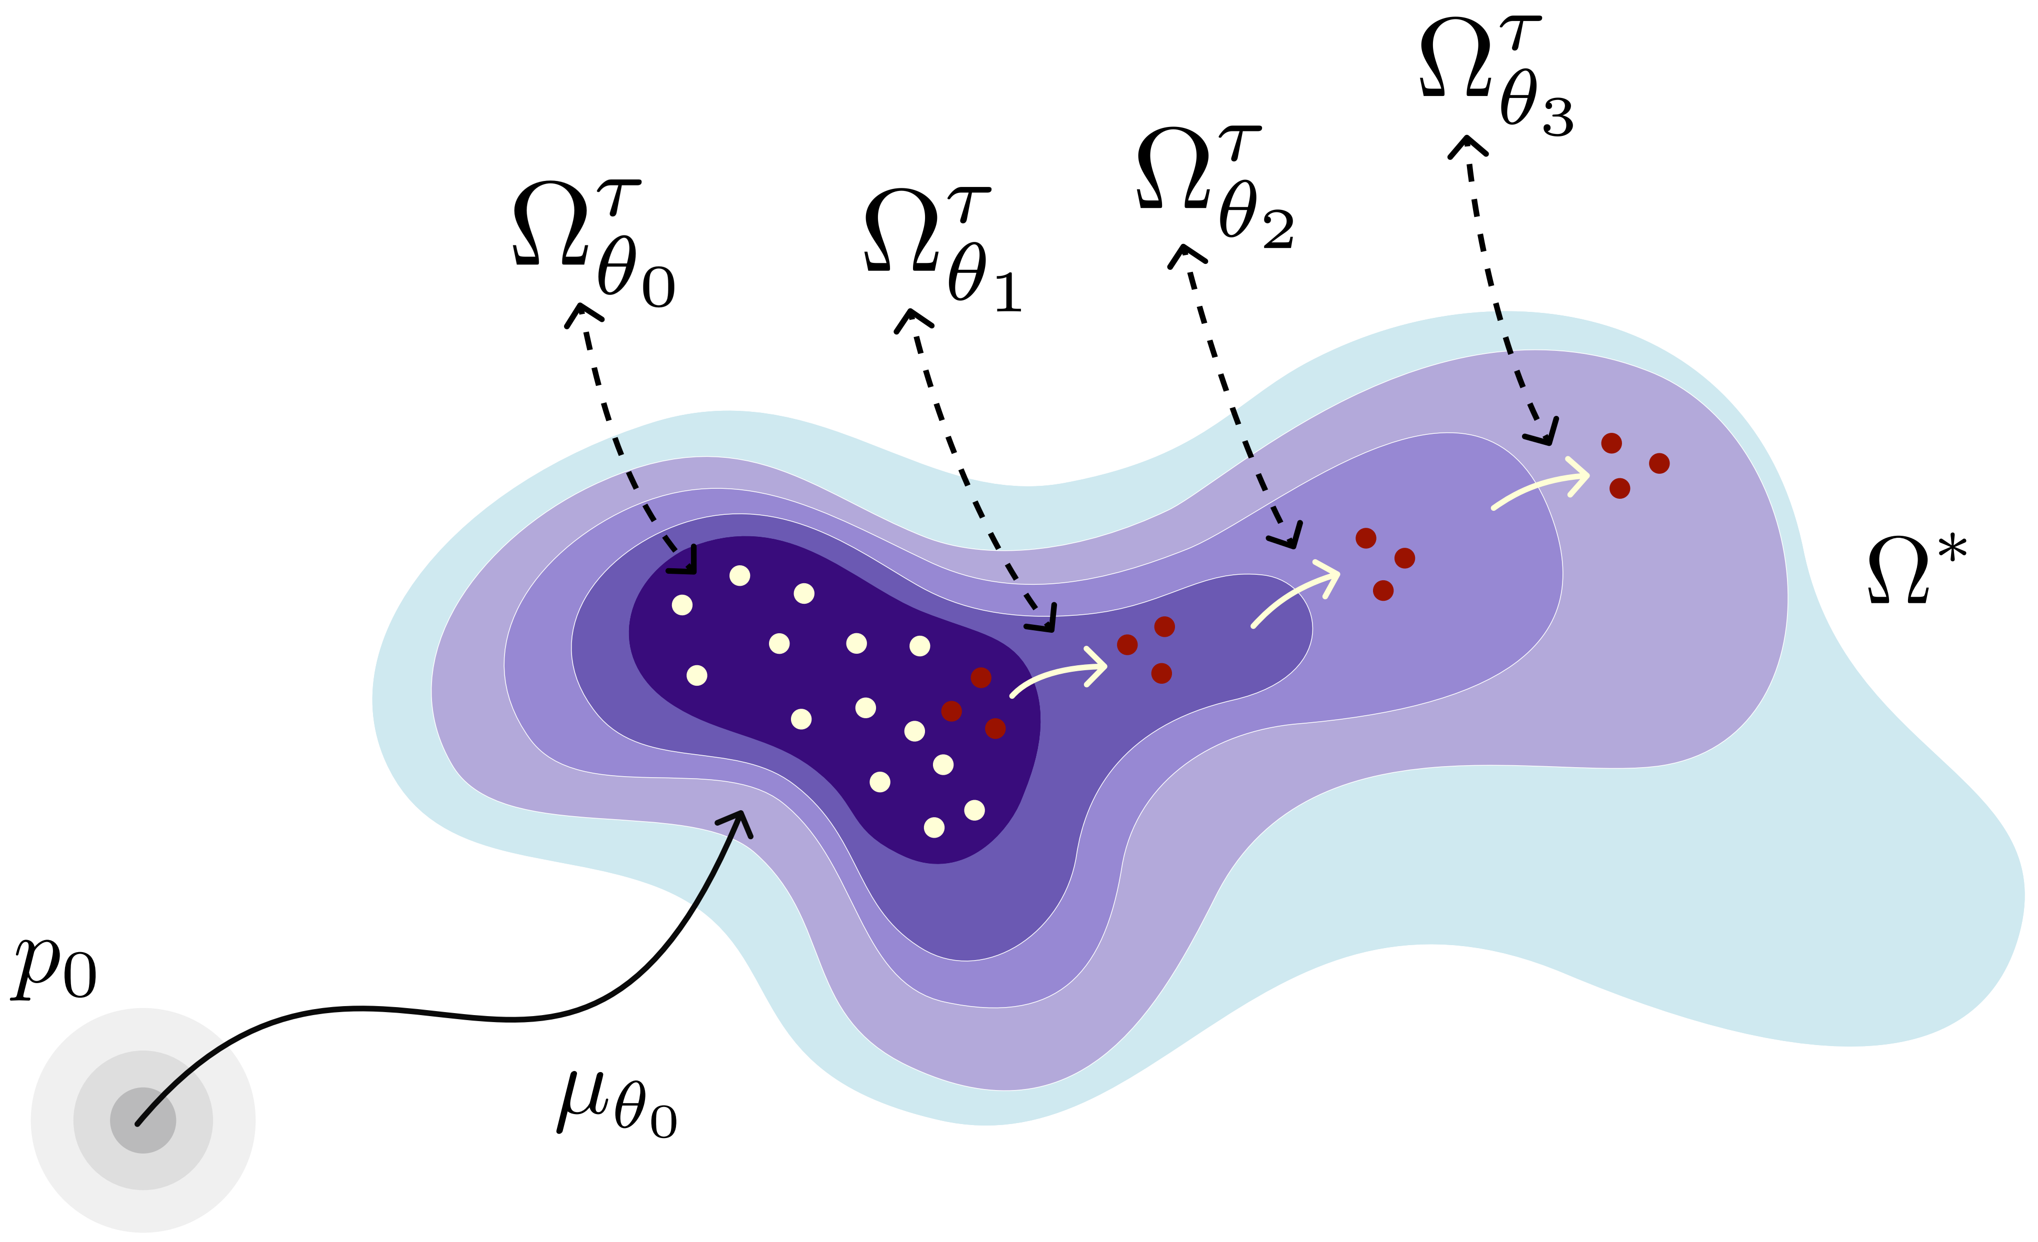

Black-box noisy binary verifier feedback

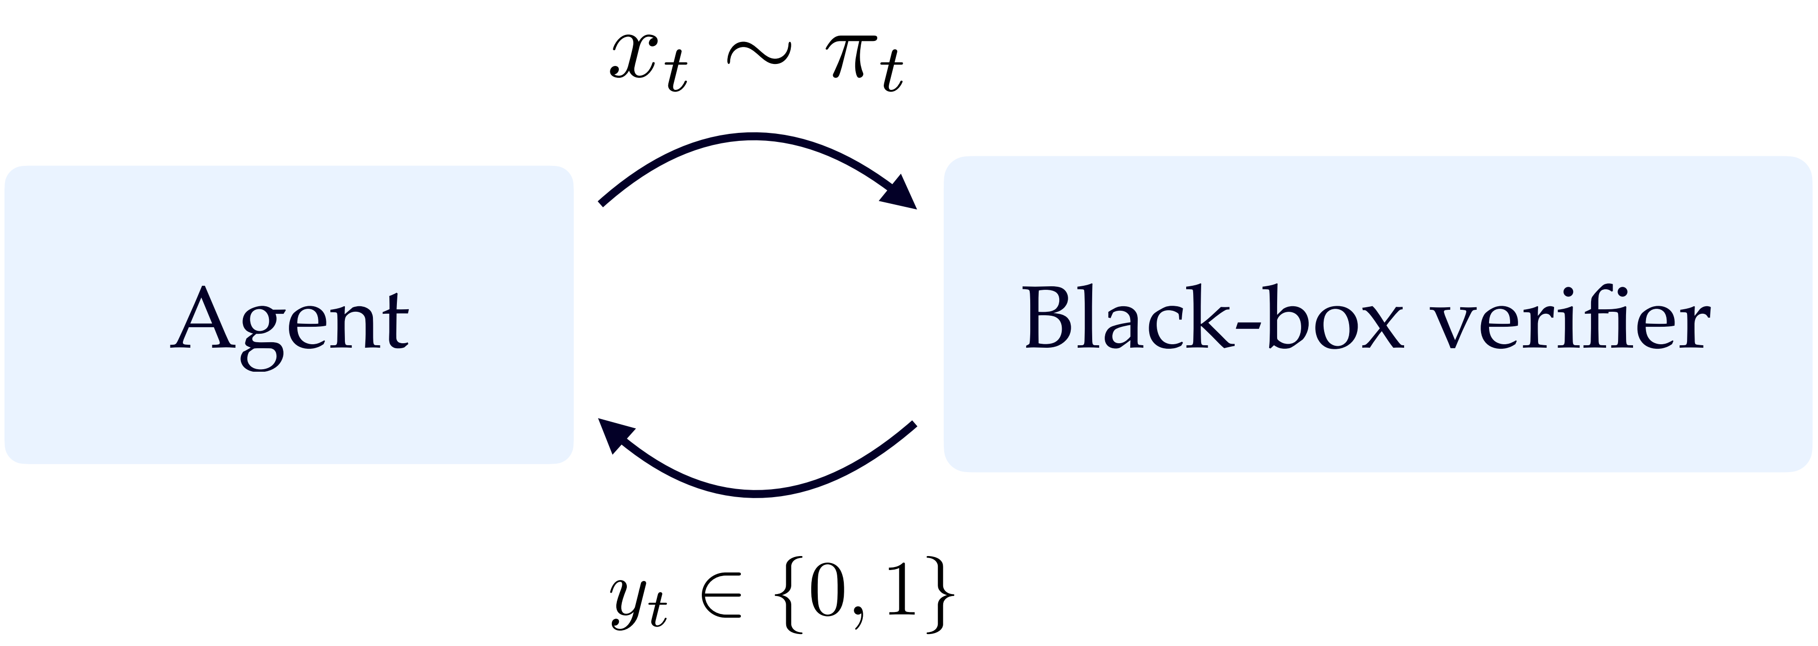

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

Generable Set Expansion for OOD Flow Modeling

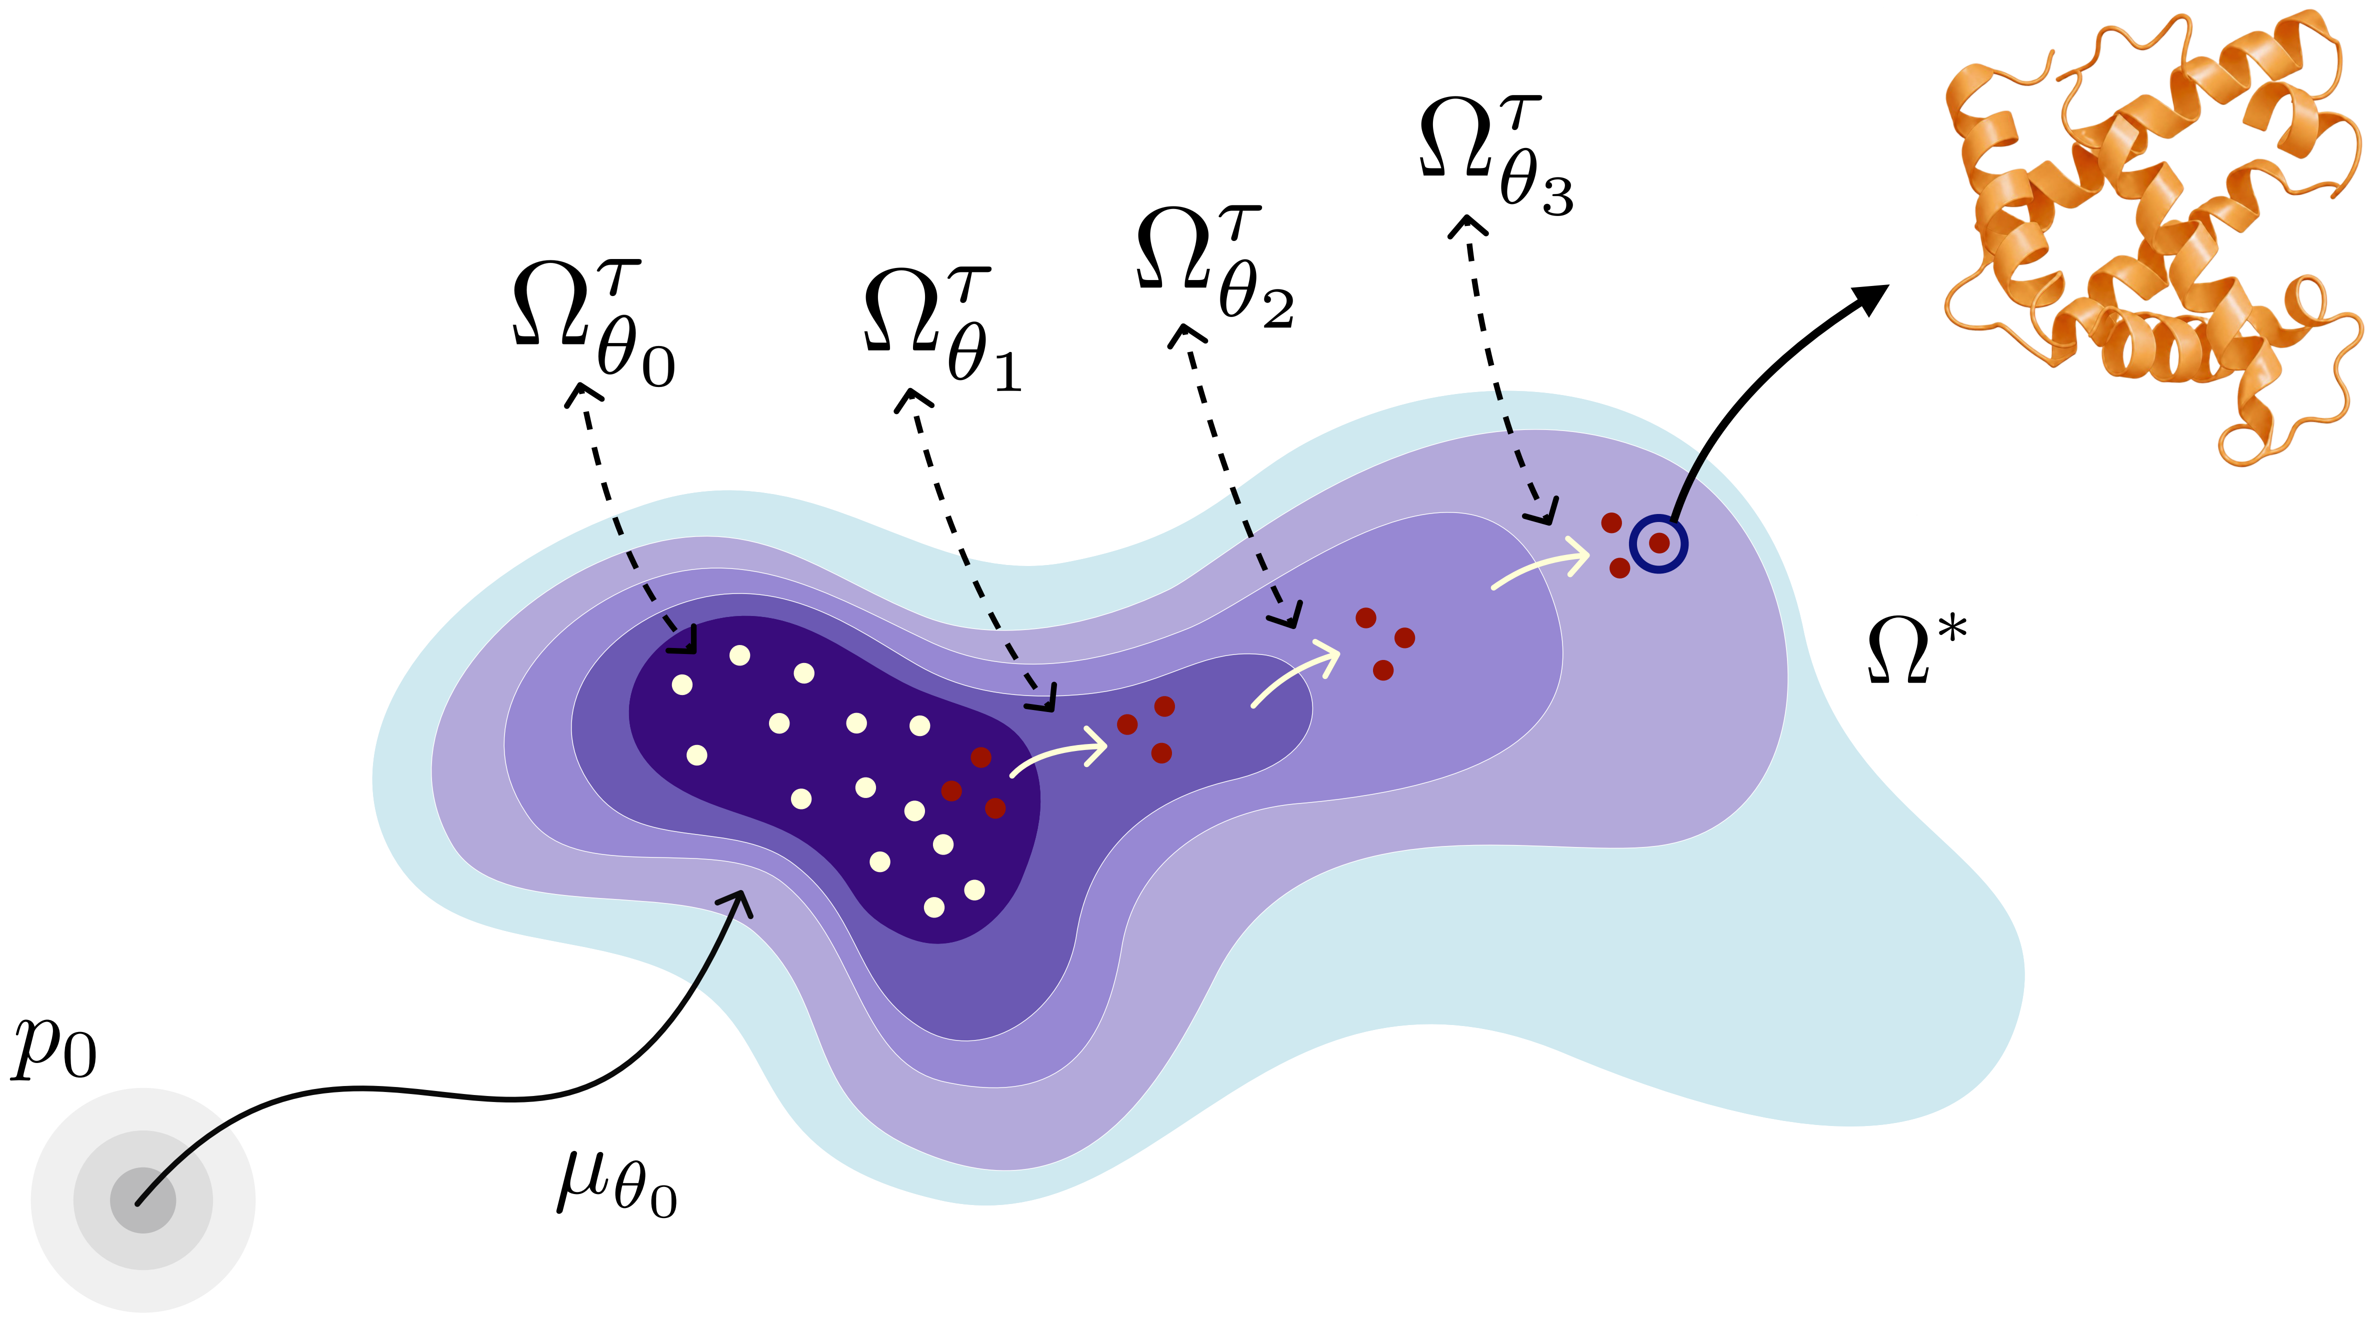

Black-box noisy binary verifier feedback

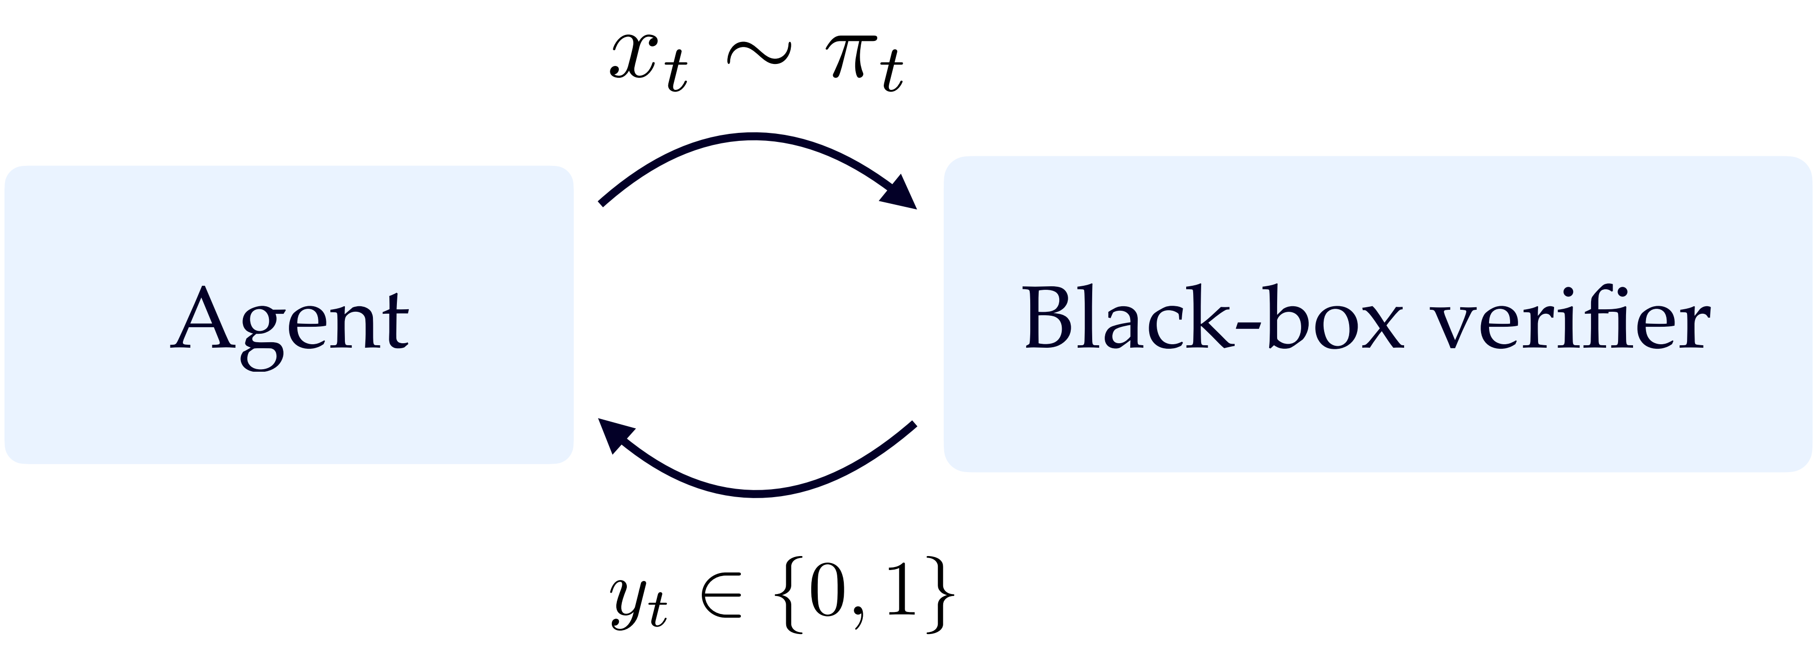

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T}$$

Generable Set Expansion for OOD Flow Modeling

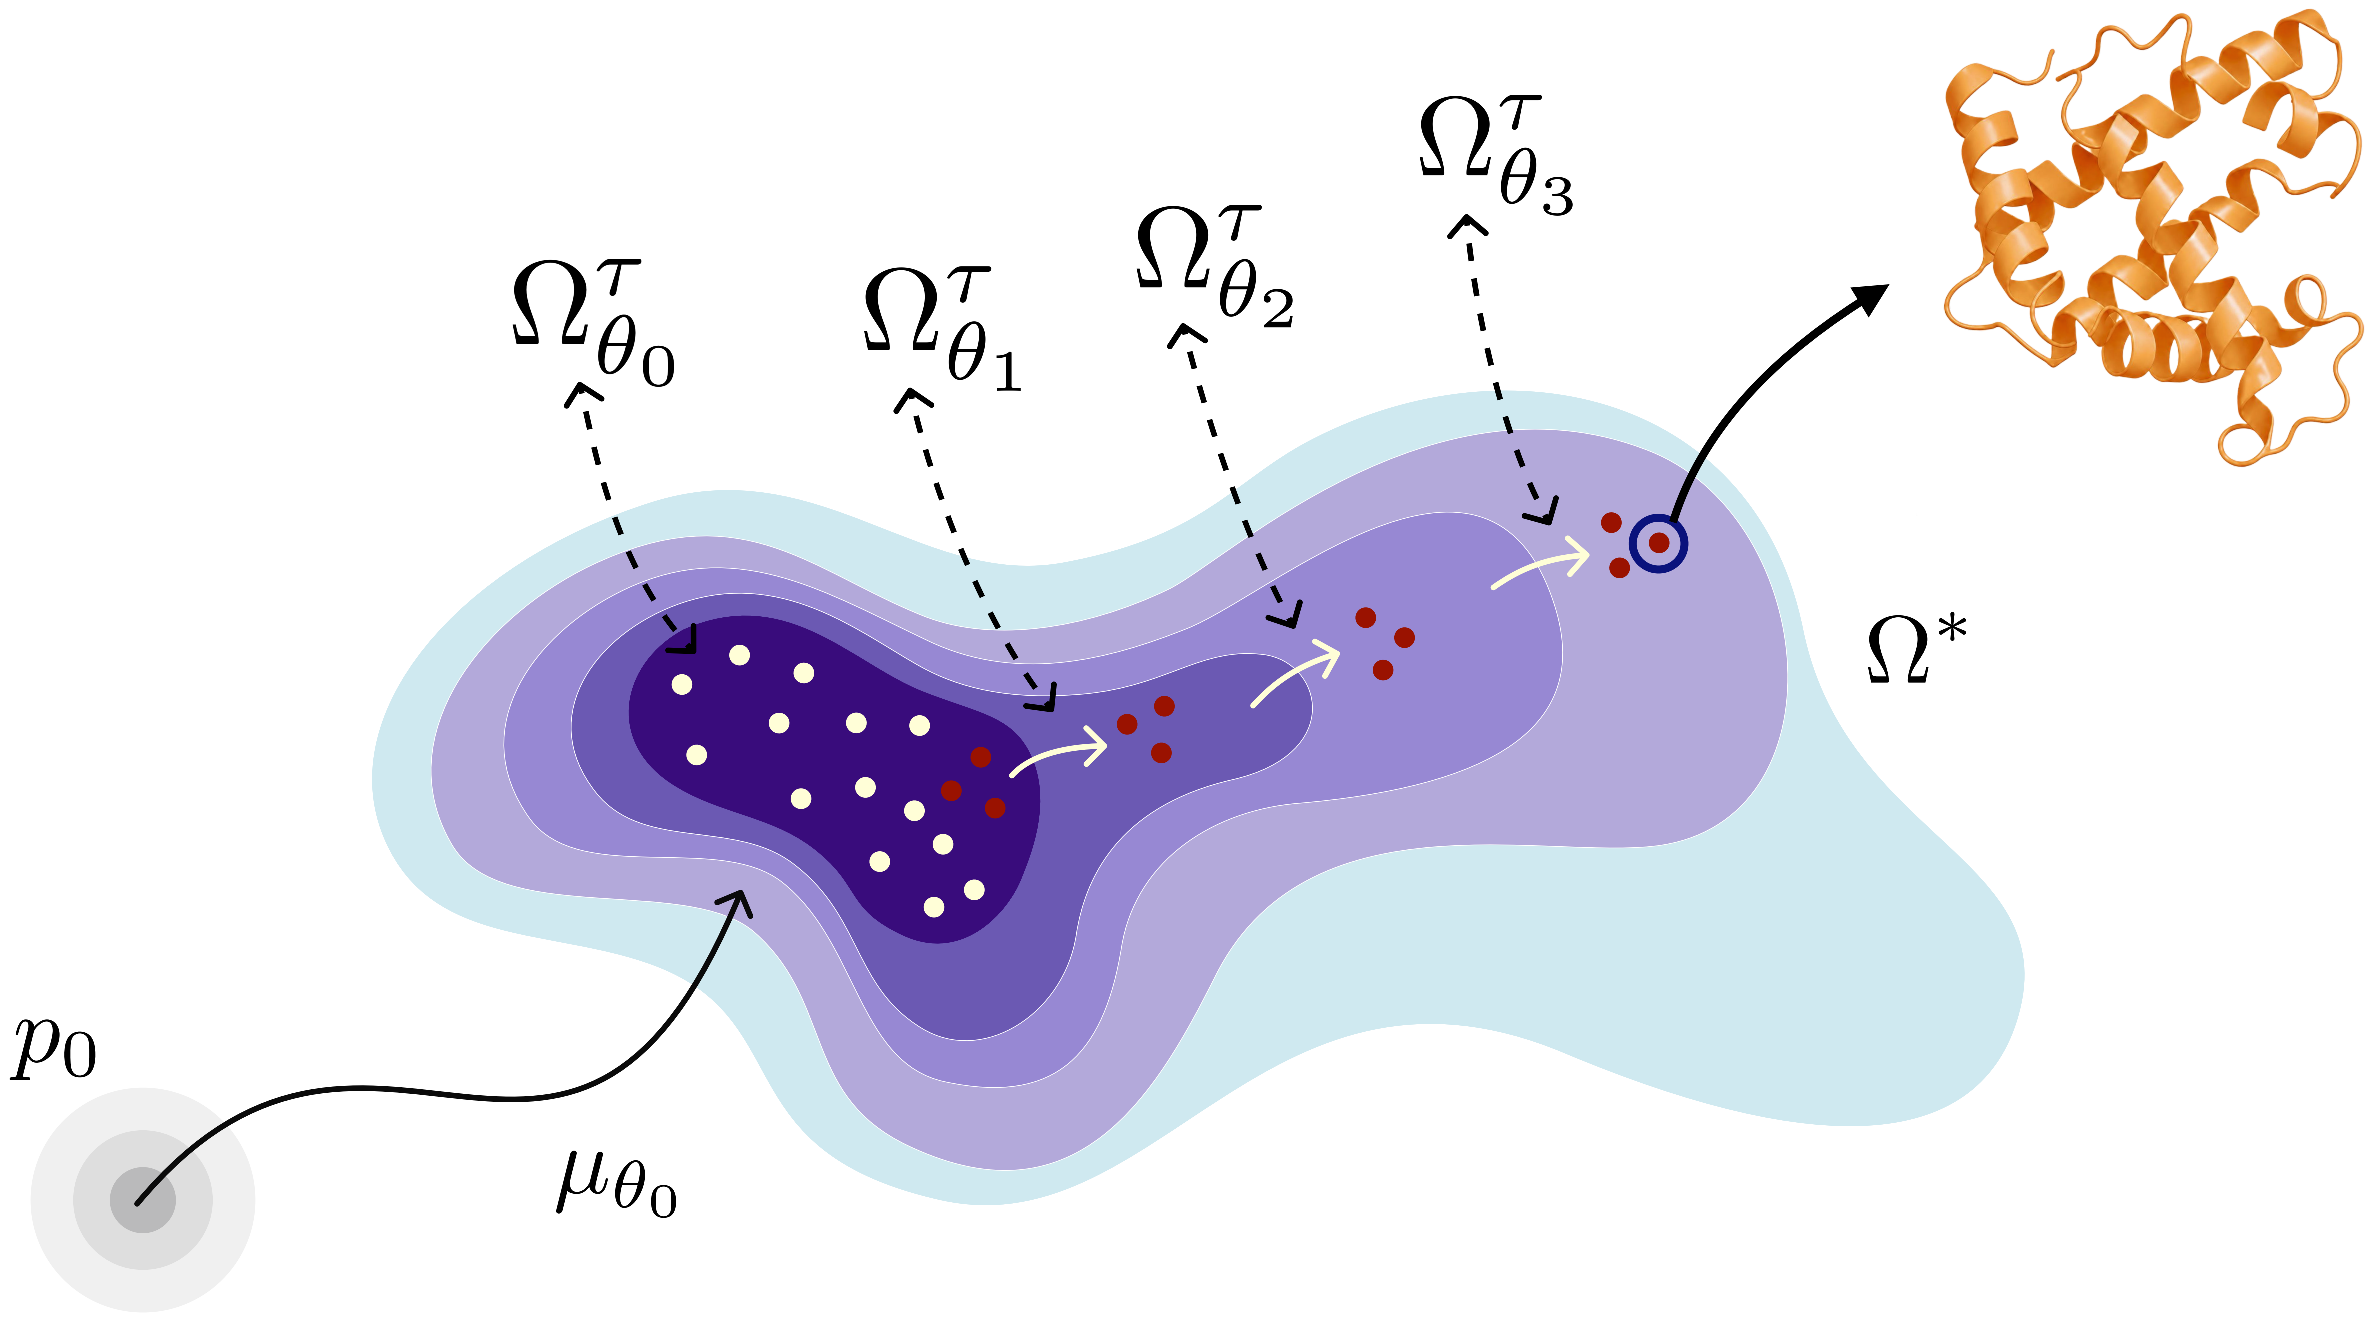

Black-box noisy binary verifier feedback

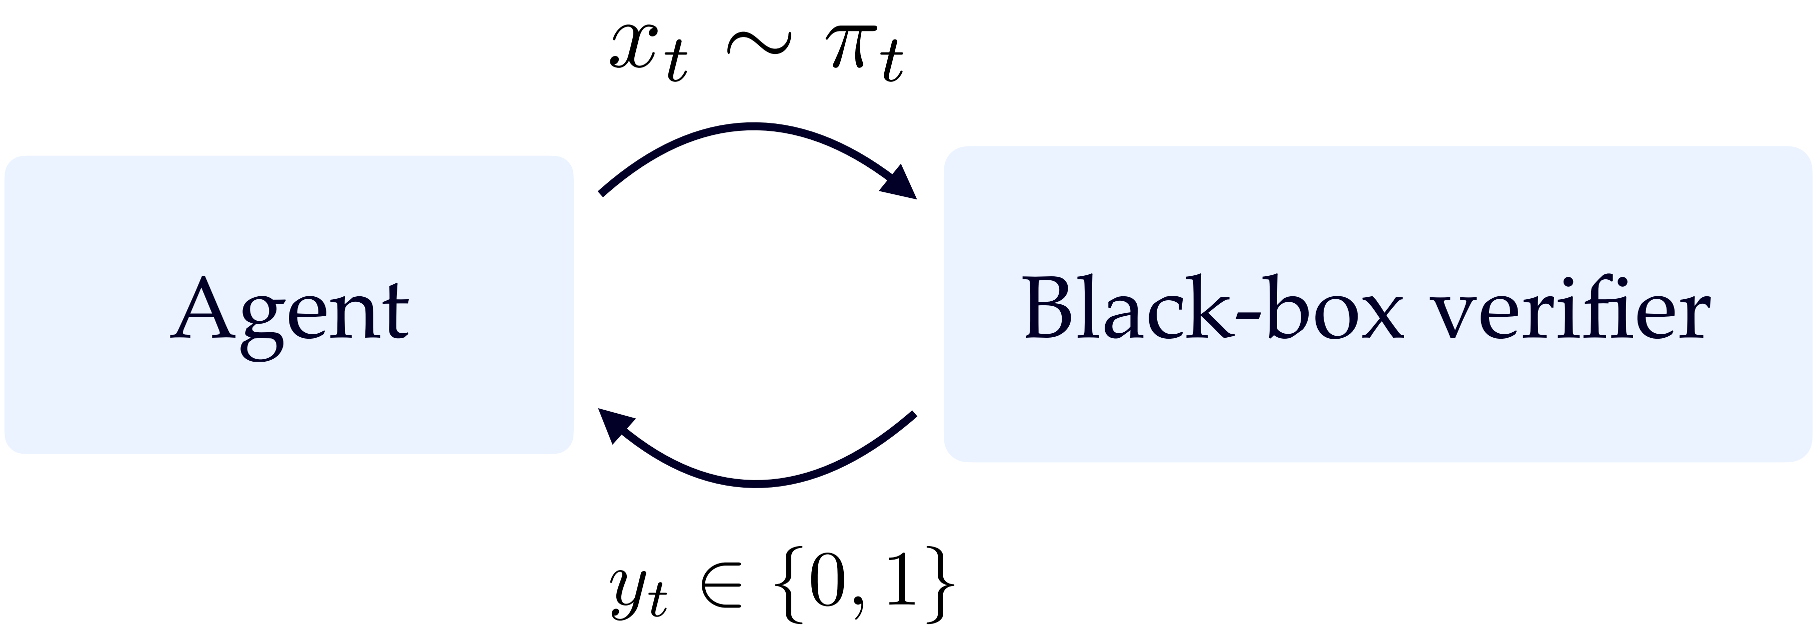

$$\theta^{\text{pre}} = \theta_0 \to \theta_1 \to \cdots \to \theta_T, \qquad \Omega^{\tau}_{\theta_0} \subseteq \Omega^{\tau}_{\theta_1} \subseteq \cdots \subseteq \Omega^{\tau}_{\theta_T} \overset{?}{\approx} \Omega^{\star}$$

How to self-generate data for expansion? Key Challenge

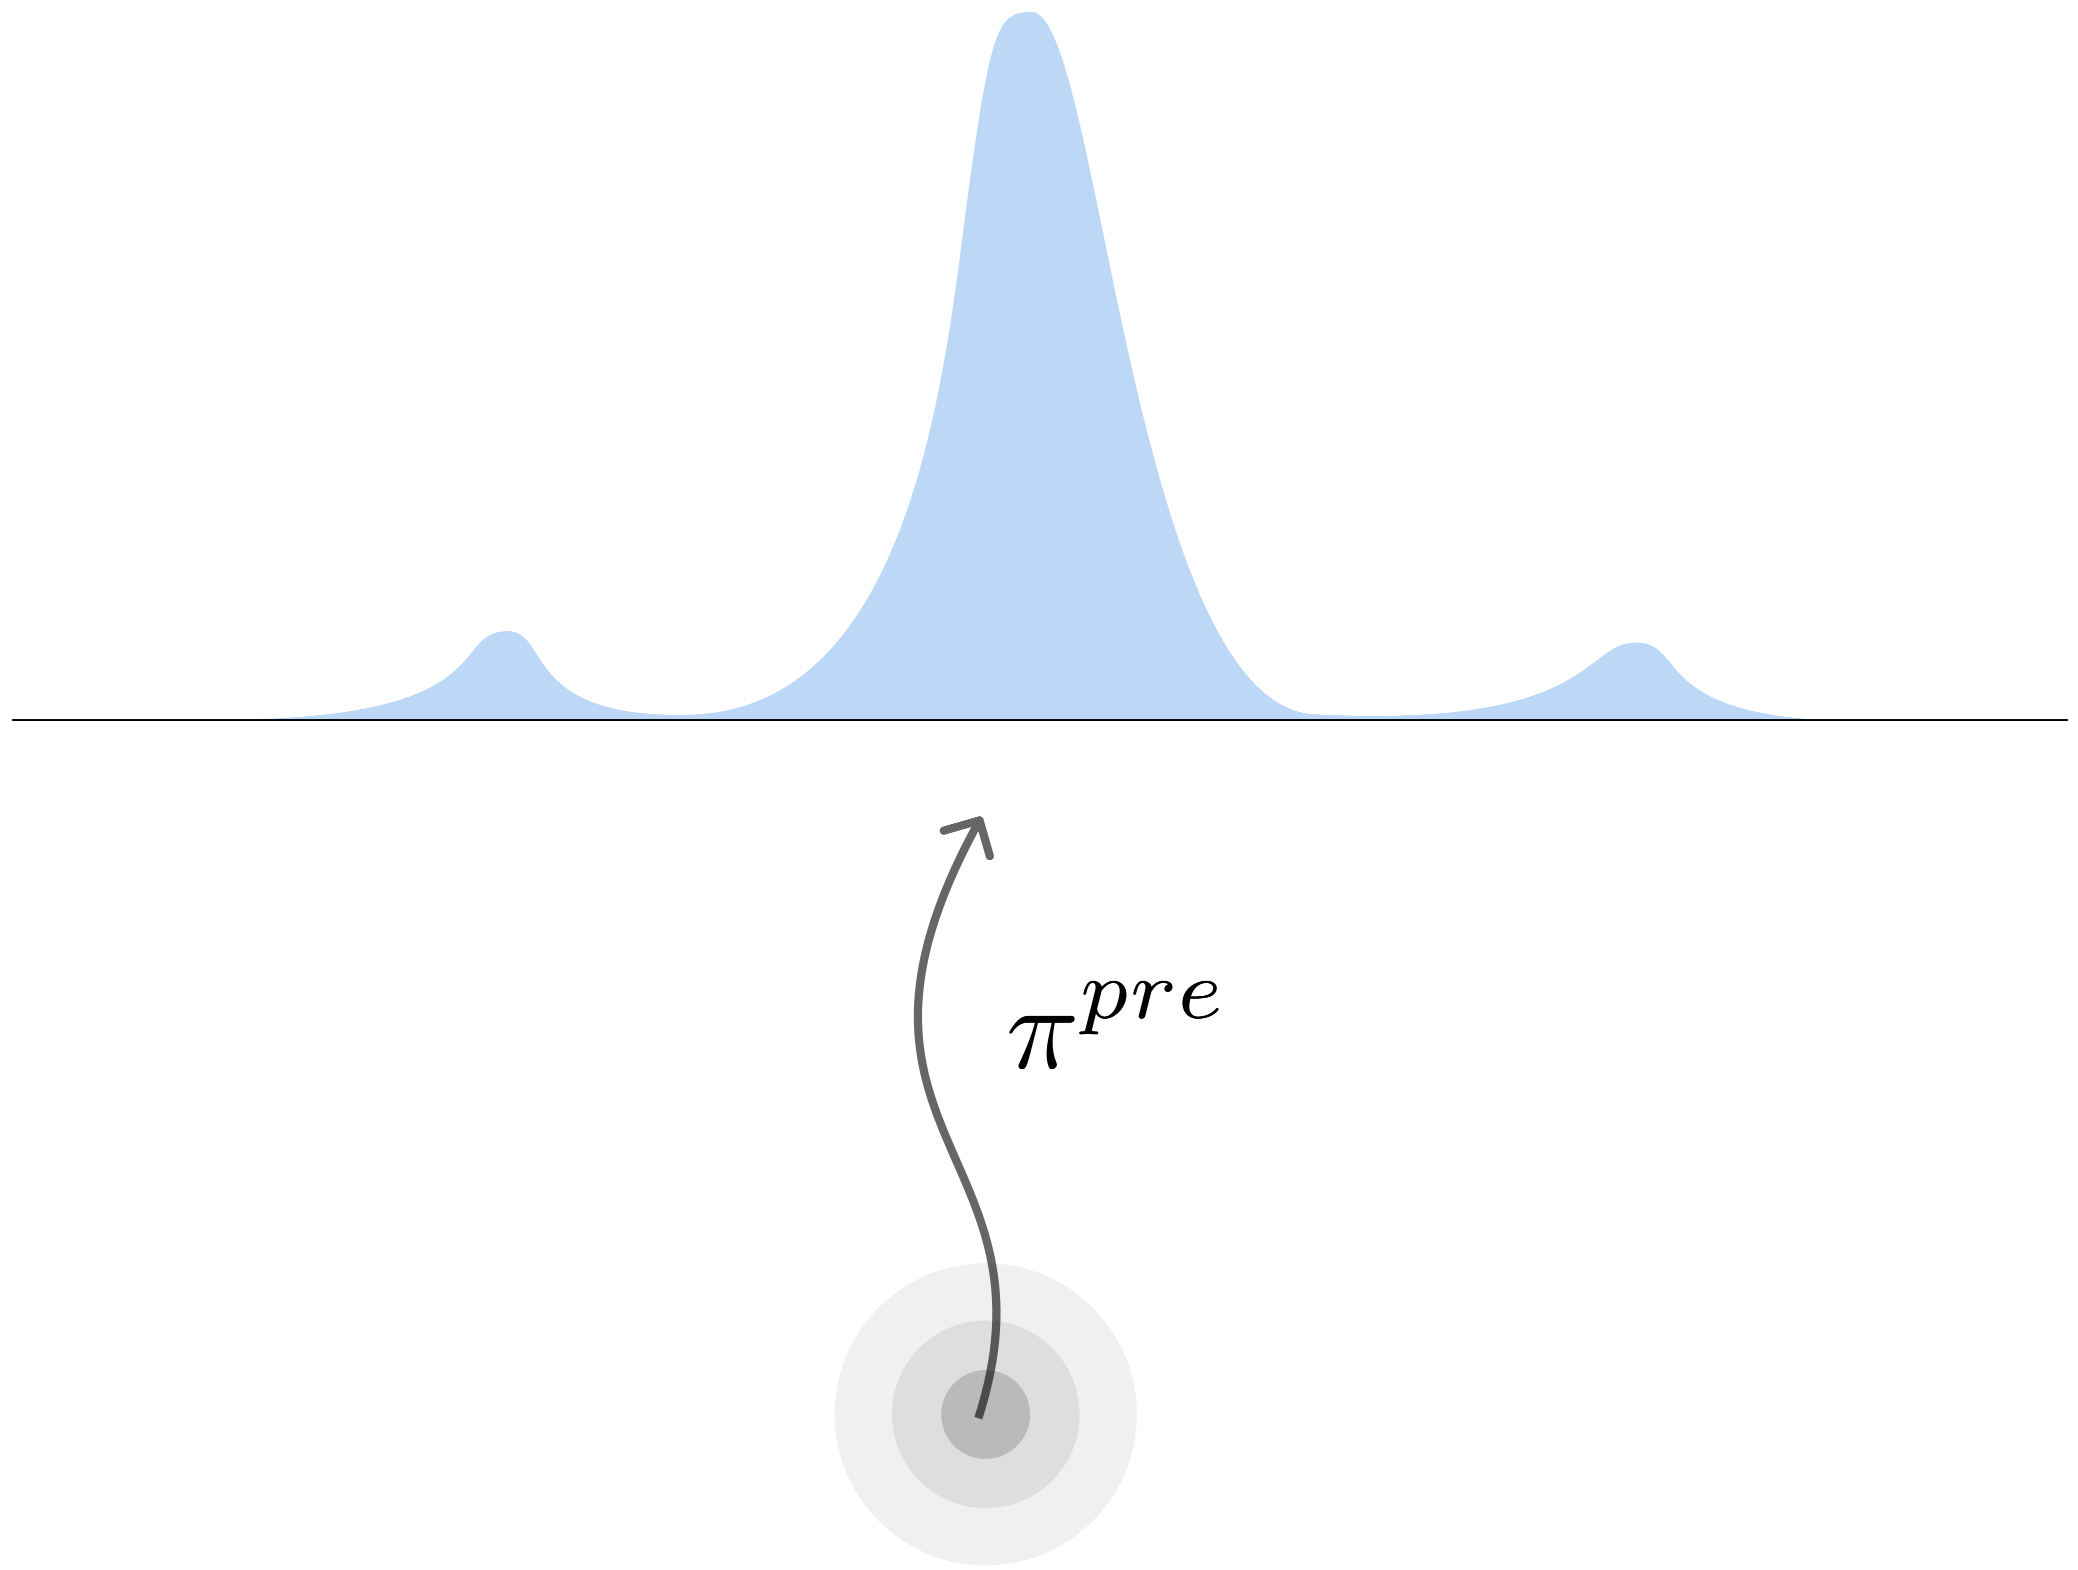

How to self-generate data for expansion? Key Challenge

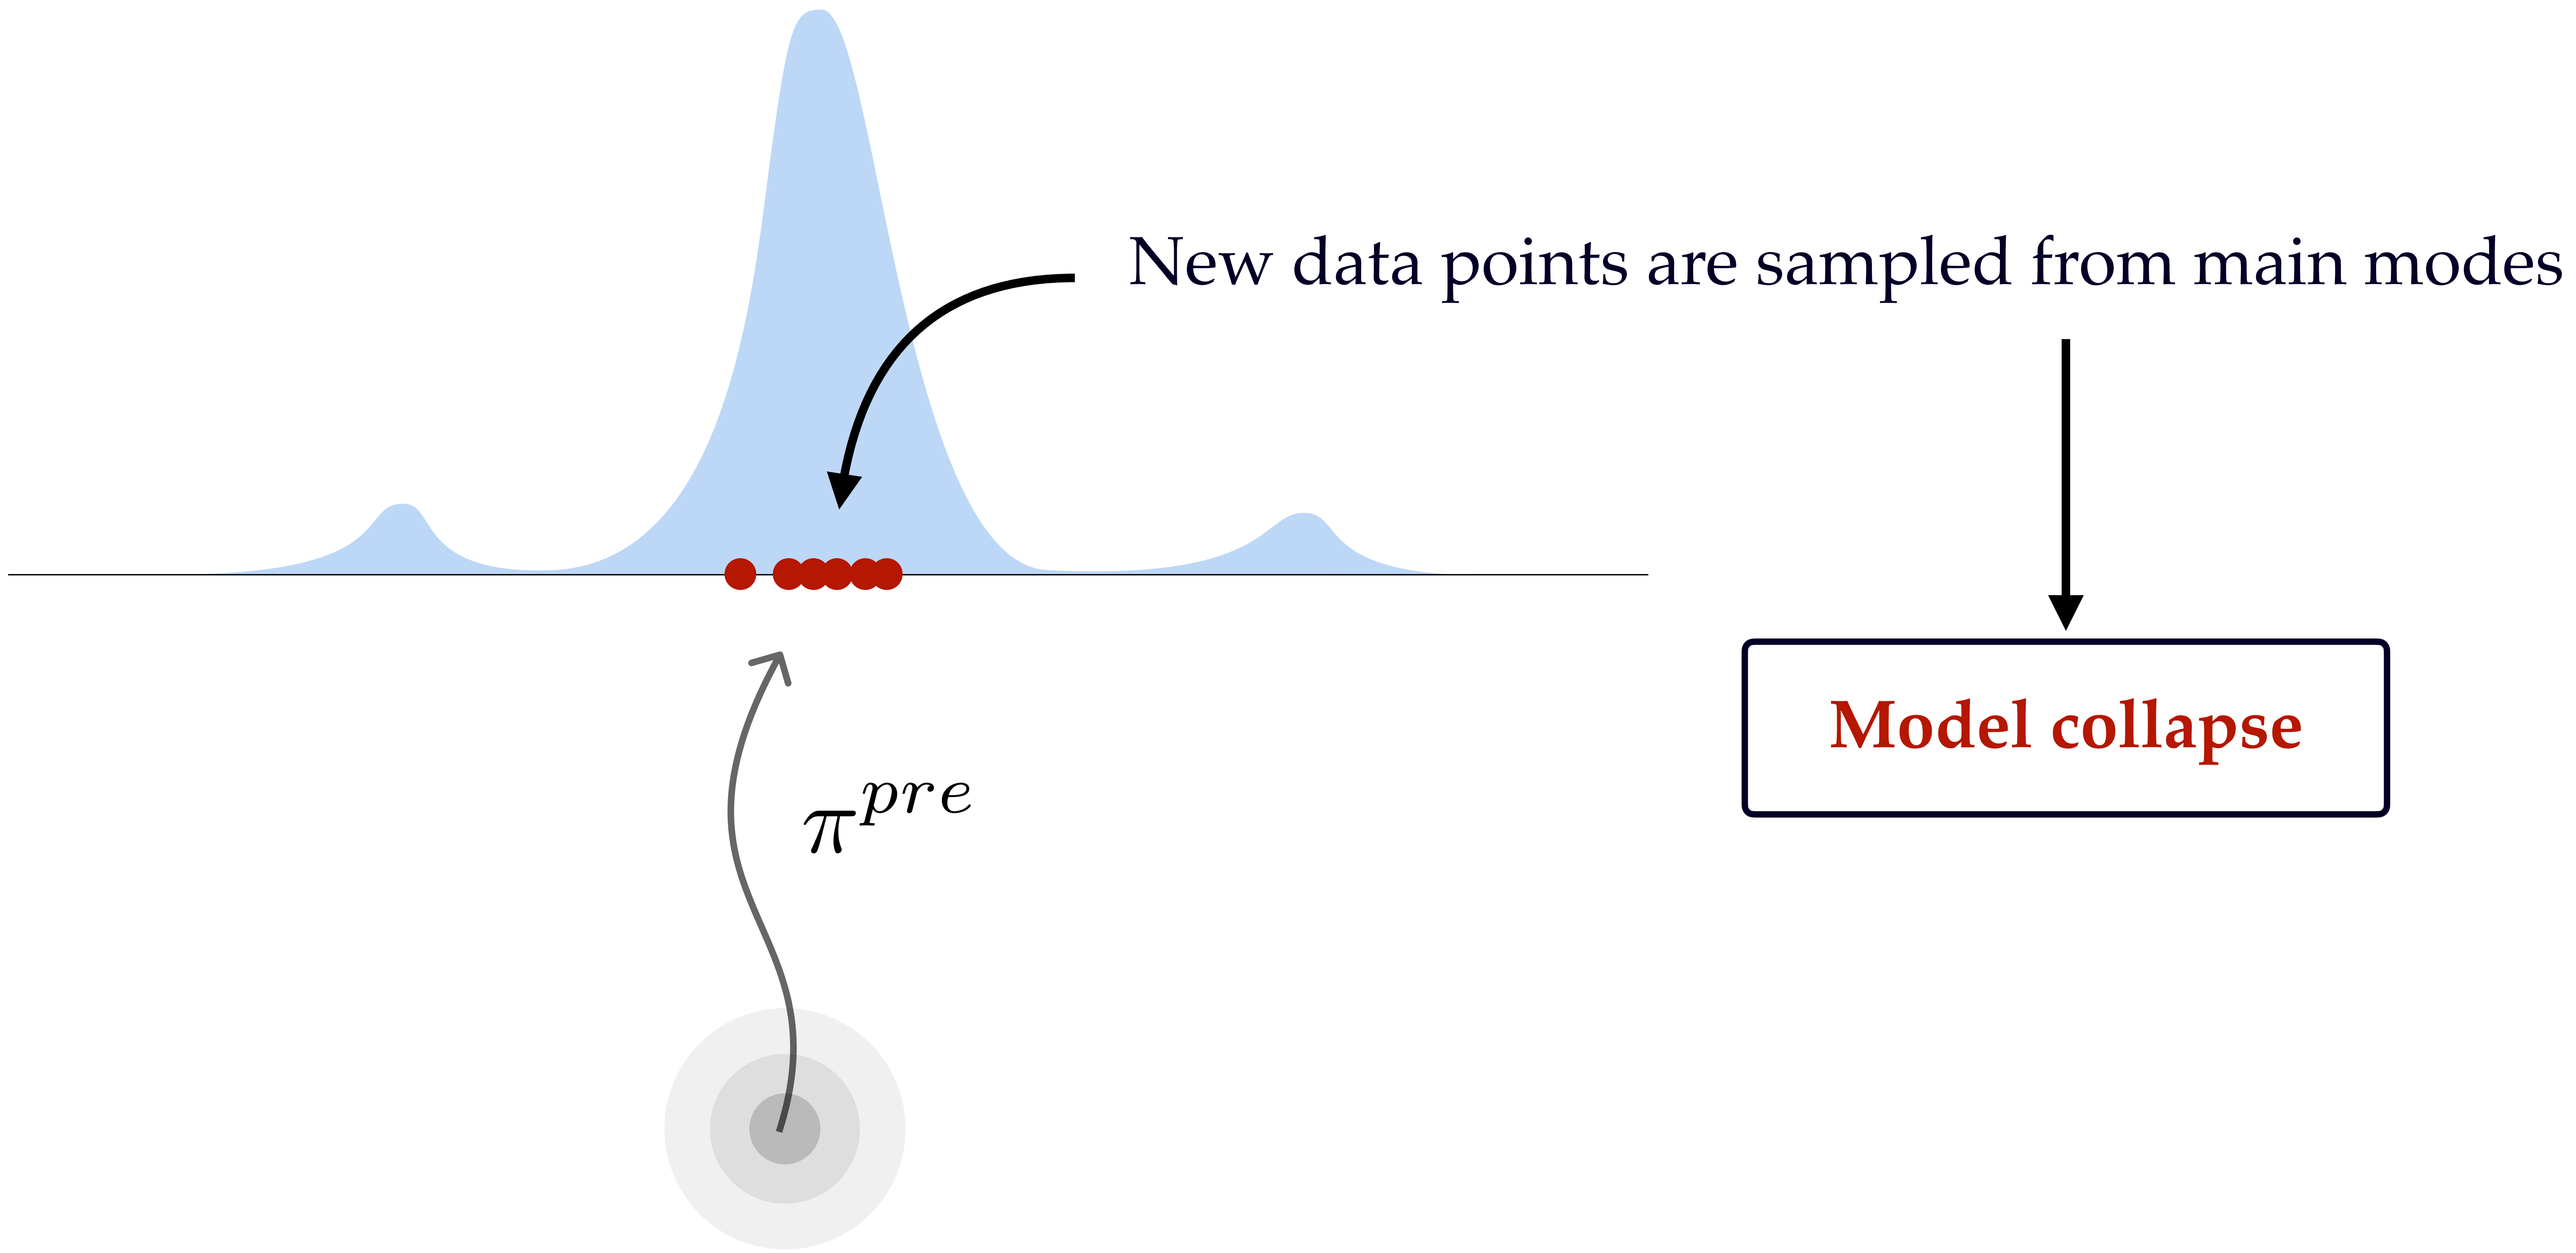

Model collapse

How to self-generate data for expansion? Generative Uncertainty Sampling

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

How to self-generate data for expansion? Generative Uncertainty Sampling

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

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

How to self-generate data for expansion? Generative Uncertainty Sampling

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

How to self-generate data for expansion? Generative Uncertainty Sampling

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

- Leverages flow-learned representation for efficient active learning
- Trades-off high uncertainty and prior signal

Active Flow Expansion (ActFlow)

Init: Initial flow model μ_{θ_0}, black-box verifier $\tilde{v}$, iterations T, $\mathcal{D}_0 \leftarrow \emptyset$

Active Flow Expansion (ActFlow)

Init: Initial flow model μ_{θ_0}, black-box verifier $\tilde{v}$, iterations T, $\mathcal{D}_0 \leftarrow \emptyset$

For $t = 0, 1, \dots, T-1$:

Update surrogate uncertainty σ_t from $\mathcal{D}_t$

Self-generate:

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

Active Flow Expansion (ActFlow)

Init: Initial flow model μ_{θ_0}, black-box verifier $\tilde{v}$, iterations T, $\mathcal{D}_0 \leftarrow \emptyset$

For $t = 0, 1, \dots, T-1$:

Update surrogate uncertainty σ_t from $\mathcal{D}_t$

Self-generate:

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

Query verifier $y_{t+1} \leftarrow \tilde{v}(x_{t+1})$

$\mathcal{D}_{t+1} \leftarrow \mathcal{D}_t \cup \{(x_{t+1}, y_{t+1})\}$

$\theta_{t+1} \leftarrow \textsc{UpdateFlow}(\theta_t, \mathcal{D}_{t+1})$

Return μ_{θ_T}

Theoretical Guarantees for Out-of-Distribution Flow Modeling

Key (Informal) Assumptions.

- **Calibrated verifier uncertainty model**
- **Generative uncertainty sampling oracle**
- **EBM update abstraction**

approximate and **local** acquisition function maximization

Assumption (Generative uncertainty sampling oracle). At round t generative uncertainty sampling returns x_t such that $z_t = \phi(x_t)$:

$$\sigma_t(z_t) \;\geq\; \frac{1}{\alpha} \max_{z \in \Omega_t^\tau} \sigma_t(z), \qquad \alpha \geq 1.$$

current latent generable set

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

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

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

$$R_\epsilon^H(S_0) \subseteq \Omega_{T^\star}^\tau$$

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

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

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

$$R_\epsilon^H(S_0) \subseteq \Omega_{T^\star}^\tau$$

maximum information gain

$$\gamma_t^A \triangleq \max_{z_1, \dots, z_t \in A} \frac{1}{2} \log \det \left(I_t + (\lambda \kappa)^{-1} K_t \right)$$

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

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

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

$$R_\epsilon^H(S_0) \subseteq \Omega_{T^\star}^\tau$$

maximum information gain

$$\gamma_t^A \triangleq \max_{z_1, \ldots, z_t \in A} \frac{1}{2} \log \det \left(I_t + (\lambda \kappa)^{-1} K_t \right)$$

generable representation (i.e., latent) set

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:

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

maximum information gain

$$\gamma^A_t \triangleq \max_{z_1,\ldots,z_t \in A} \frac{1}{2} \log\det\left(I_t + (\lambda\kappa)^{-1} K_t\right)$$

generable representation (i.e., latent) set

$R^H_\epsilon(S_0)$ **H-fold** recursive application of **one-step reachability operator** over learned flow representation

$$R_\epsilon(S) := \{z \in \mathcal{Z} : \exists z' \in S \text{ s.t. } s(g(z')) - L_s L_g d(z, z') - \epsilon \geq h\}$$

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

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

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

$$R_\epsilon^H(S_0) \subseteq \Omega_{T^\star}^\tau$$

maximum information gain

$$\gamma_t^A \triangleq \max_{z_1,\ldots,z_t \in A} \frac{1}{2} \log\det\left(I_t + (\lambda\kappa)^{-1} K_t\right)$$

generable representation (i.e., latent) set

$S_0 \subset \Omega_0^\tau$ subset of pre-trained model generable set containing valid designs

$R_\epsilon^H(S_0)$ **H-fold** recursive application of **one-step reachability operator** over learned flow representation

$$R_\epsilon(S) := \{z \in \mathcal{Z} : \exists z' \in S \text{ s.t. } s(g(z')) - L_s L_g d(z, z') - \epsilon \geq h\}$$

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:

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

Allows to identify **structural conditions** and **statistical complexity** for:

$$\Omega^\star \subseteq \Omega^{X,\tau_X}_{T^\star}$$

valid design space

generable set on $\mathcal{X}$

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

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

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

$$R_{\epsilon}^{H}(S_0) \subseteq \Omega_{T^{\star}}^{\tau}$$

Takeaway

ActFlow allows to build (generable) **set-theoretic guarantees** for **out-of-distribution flow modeling**.

Distribution matching

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

→

Design space coverage

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

ActFlow for Chemical and Biological Generative Discovery

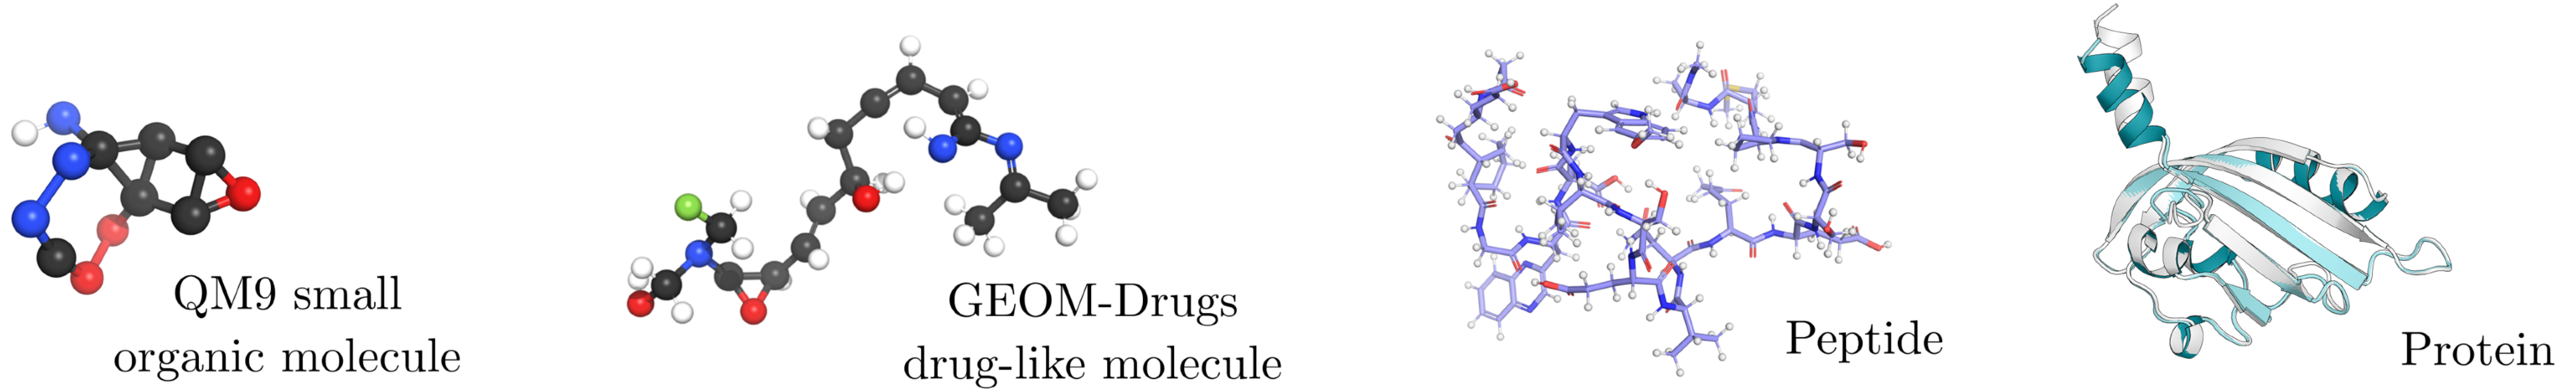

ActFlow for Small Organic and Drug-like Molecules

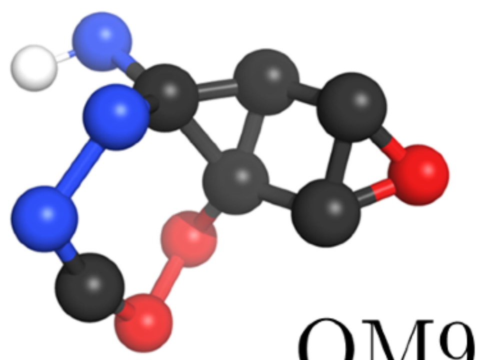

QM9 small organic molecule

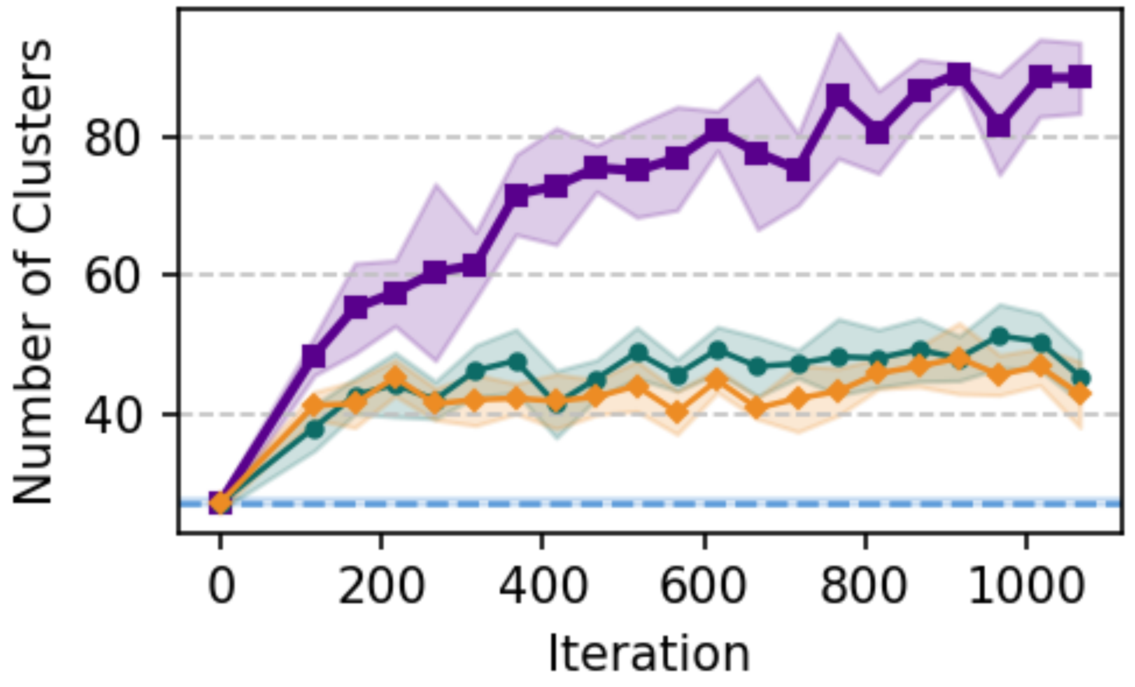

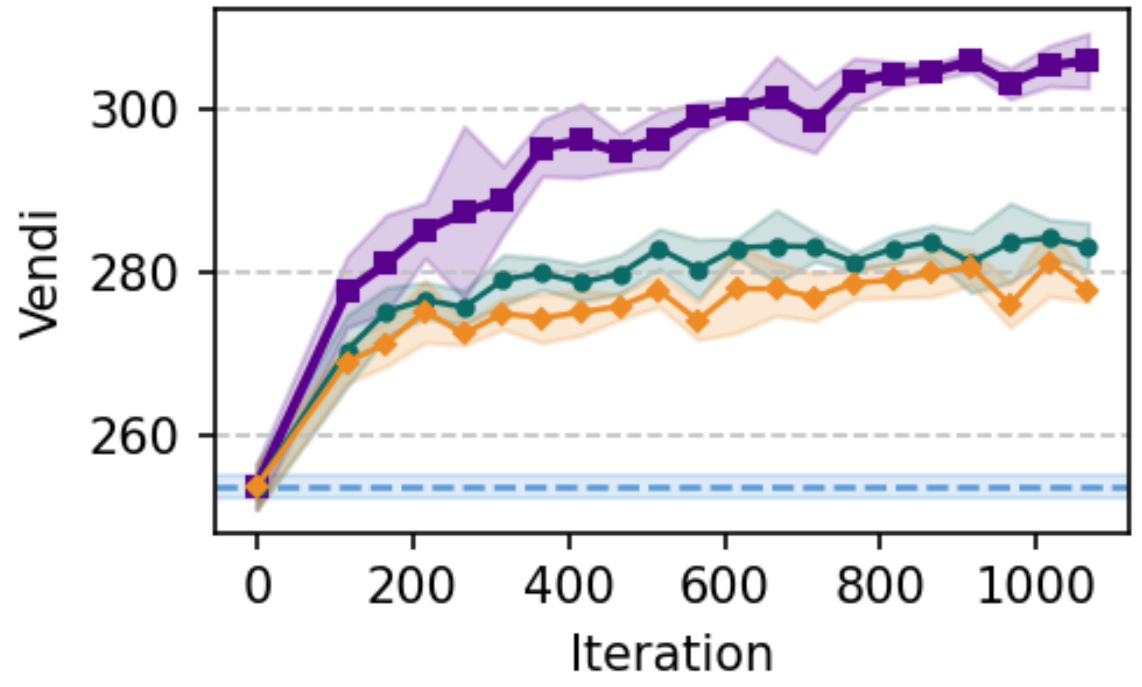

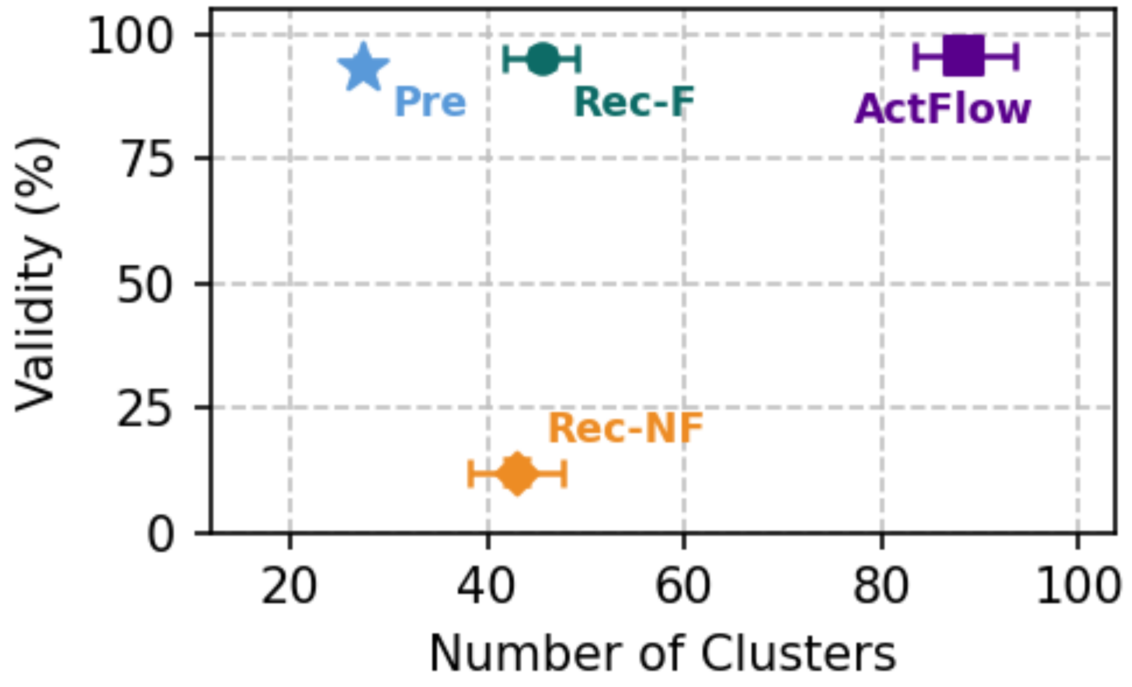

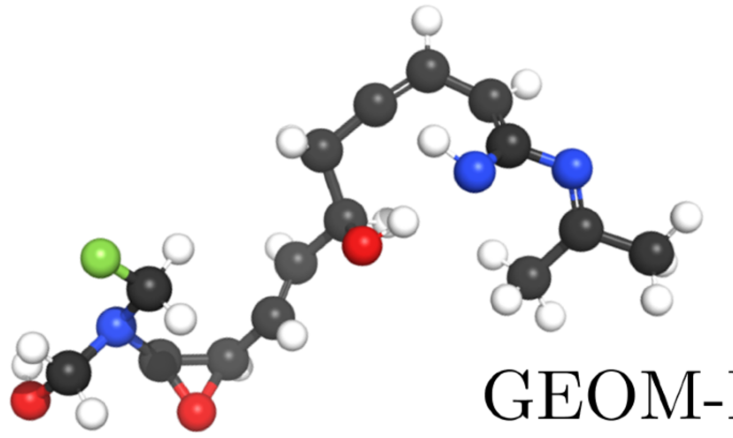

GEOM-Drugs drug-like molecule

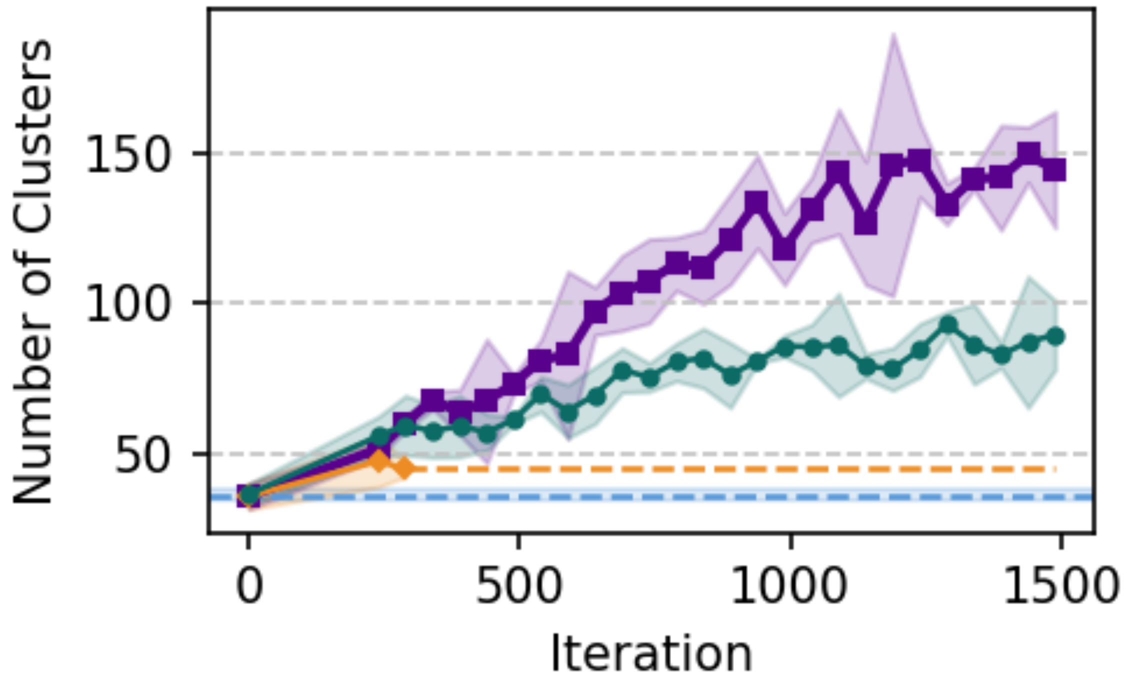

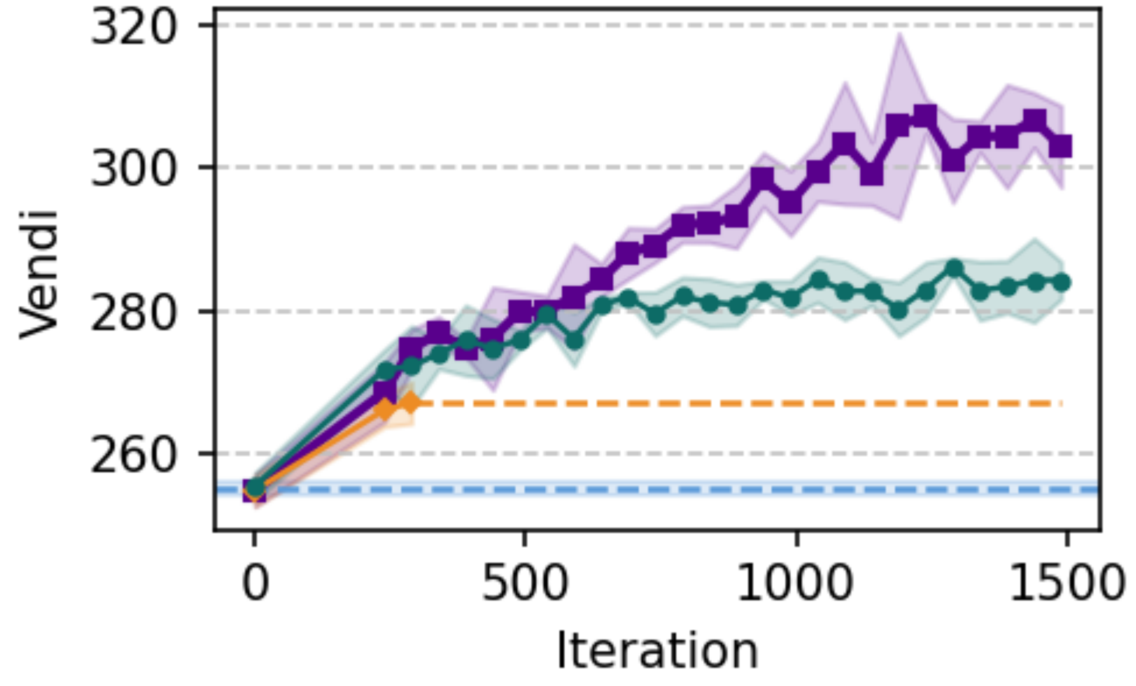

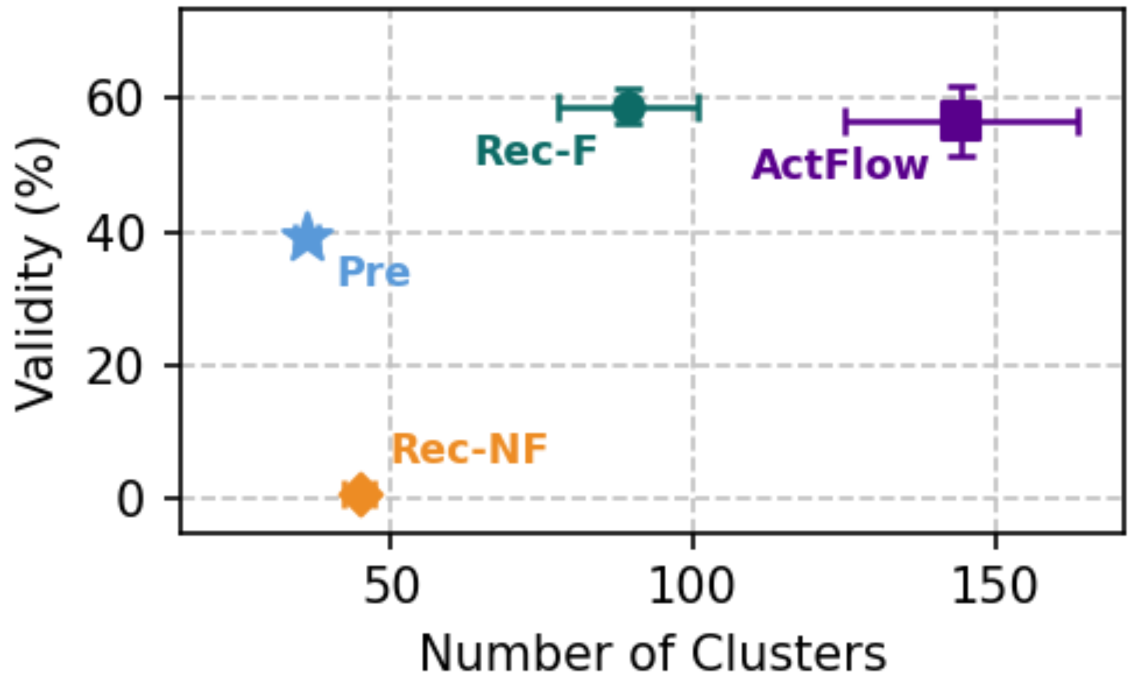

ActFlow for Therapeutic Peptides and Protein Design

	Therapeutic Peptide Design				Protein Sequence Design			
Method	**Coverage ↑**	**Diversity ↑**	**FID**	**Validity (%) ↑**	**Coverage ↑**	**Diversity ↑**	**FID**	**Validity (%) ↑**
Pre	61.00 ± 0.00	13.73 ± 0.00	0.00 ± 0.00	40.12 ± 0.00	66.50 ± 5.63	12.87 ± 0.63	0.05 ± 0.01	70.81 ± 1.12
REC-NF	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	63.75 ± 11.45	11.85 ± 0.34	0.27 ± 0.06	67.81 ± 3.11
REC-F	59.67 ± 17.97	13.62 ± 4.34	3.18 ± 1.77	71.16 ± 5.22	49.50 ± 9.60	11.67 ± 0.40	0.25 ± 0.04	88.12 ± 1.42
ACTFLOW	358.33 ± 95.45	58.87 ± 25.98	59.15 ± 44.31	41.59 ± 10.06	102.75 ± 18.36	42.14 ± 10.85	5.45 ± 3.34	83.74 ± 2.70

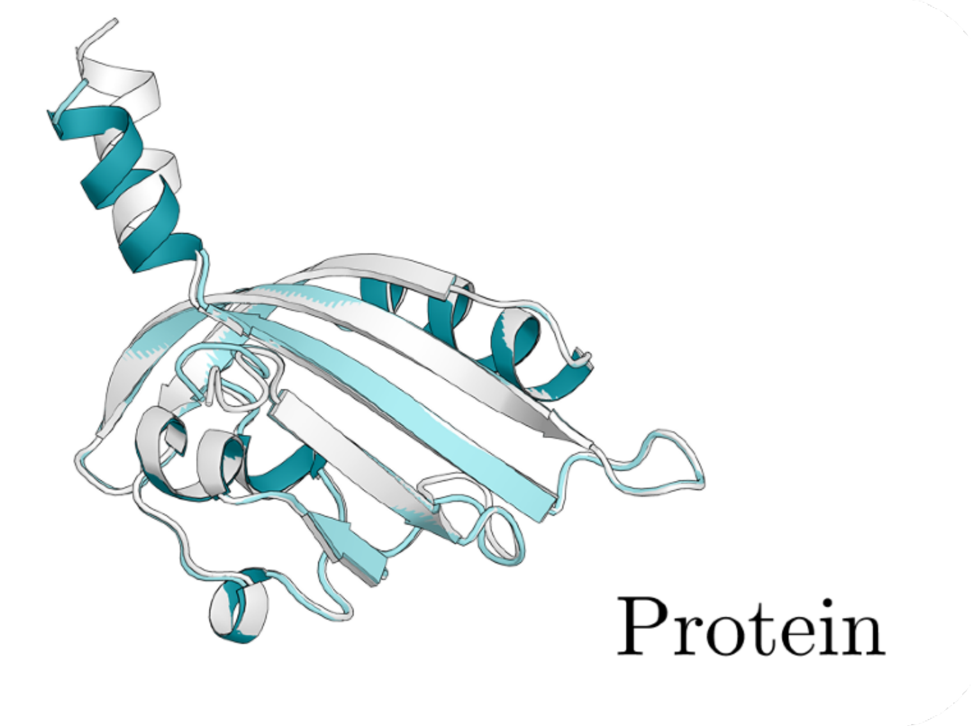

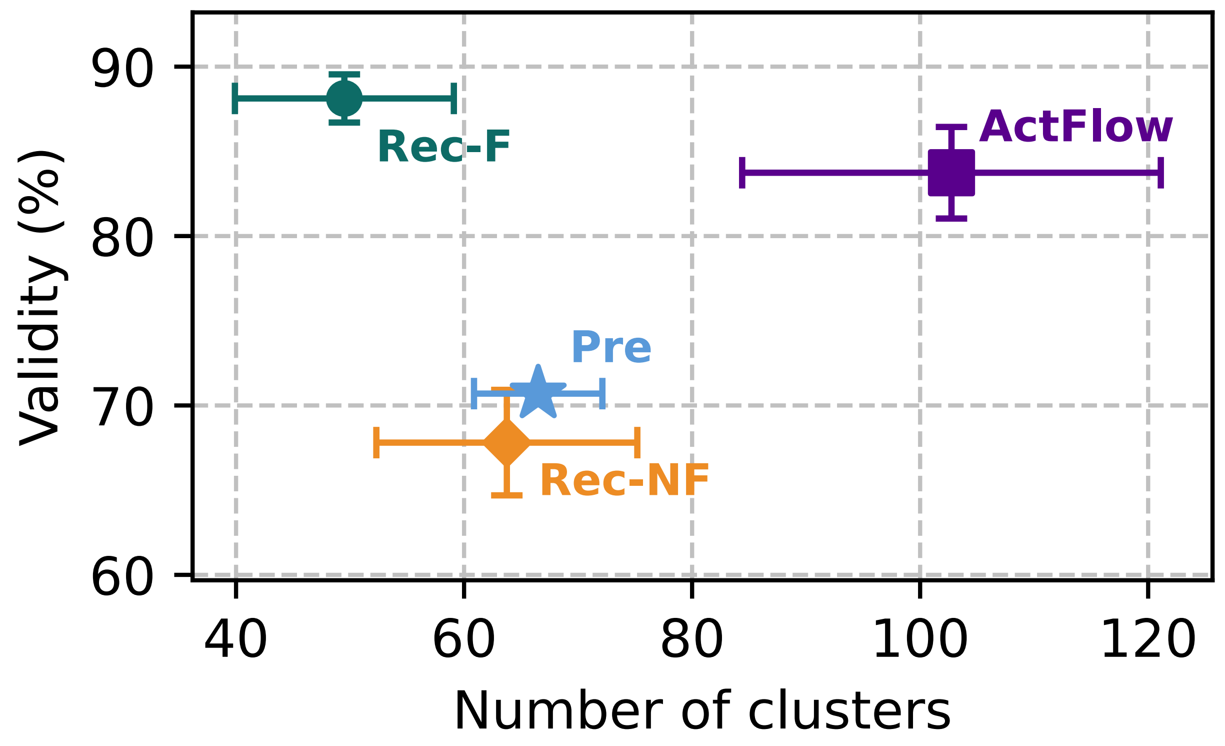

Takeaway

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

ActFlow for Therapeutic Peptides and Protein Design

Method	Therapeutic Peptide Design				Protein Sequence Design			
	Coverage ↑	Diversity ↑	FID	Validity (%) ↑	Coverage ↑	Diversity ↑	FID	Validity (%) ↑
Pre	61.00 ± 0.00	13.73 ± 0.00	0.00 ± 0.00	40.12 ± 0.00	66.50 ± 5.63	12.87 ± 0.63	0.05 ± 0.01	70.81 ± 1.12
REC-NF	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	63.75 ± 11.45	11.85 ± 0.34	0.27 ± 0.06	67.81 ± 3.11
REC-F	59.67 ± 17.97	13.62 ± 4.34	3.18 ± 1.77	71.16 ± 5.22	49.50 ± 9.60	11.67 ± 0.40	0.25 ± 0.04	88.12 ± 1.42
ACTFLOW	358.33 ± 95.45	58.87 ± 25.98	59.15 ± 44.31	41.59 ± 10.06	102.75 ± 18.36	42.14 ± 10.85	5.45 ± 3.34	83.74 ± 2.70

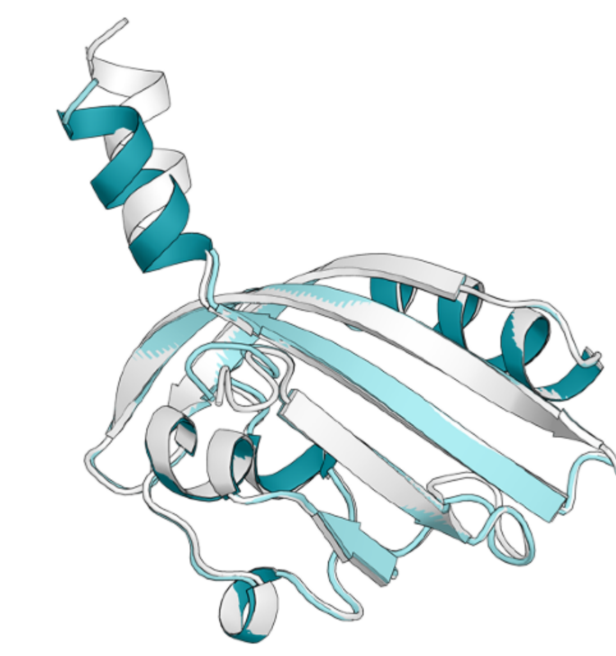

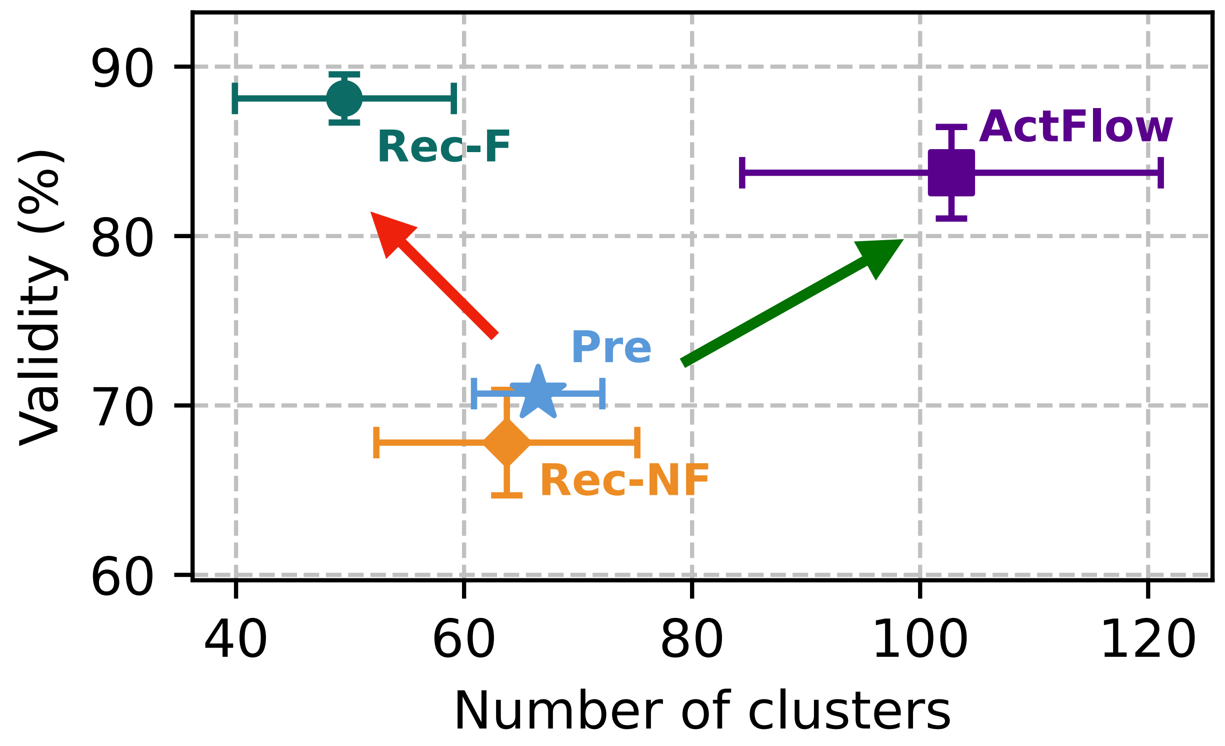

Takeaway

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

This talk:

Foundations of Generative Discovery Beyond the Data

Part I

Tail-Aware Reward Adaptation

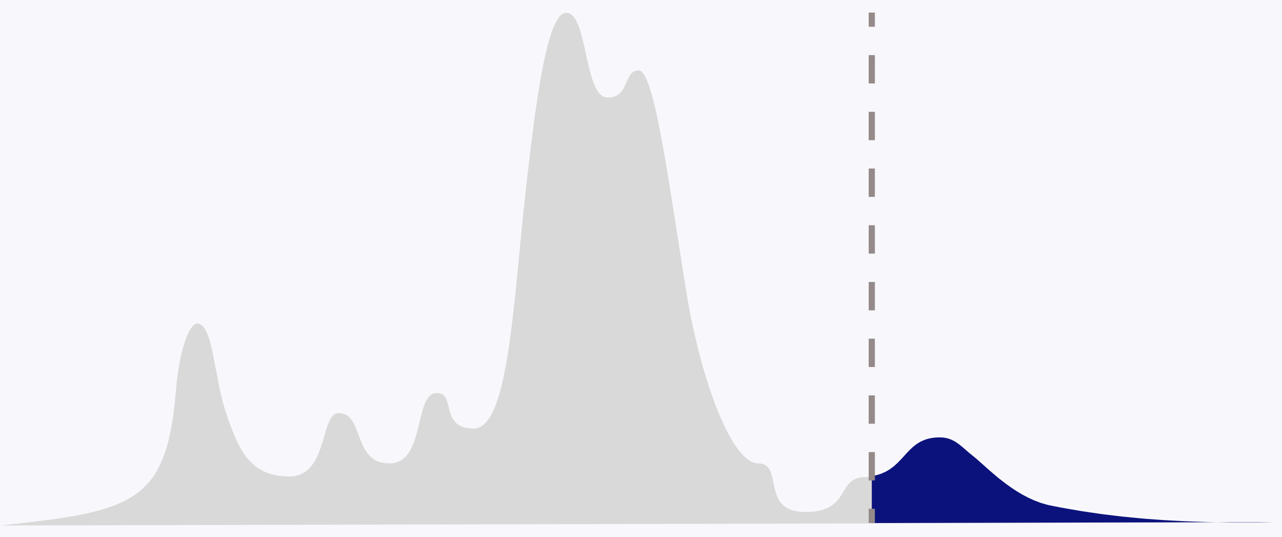

$\mathrm{CVaR}^{f}_{1-\beta}(p_1^{\pi})$

Part II

Data Debiasing and Hidden Mode Discovery

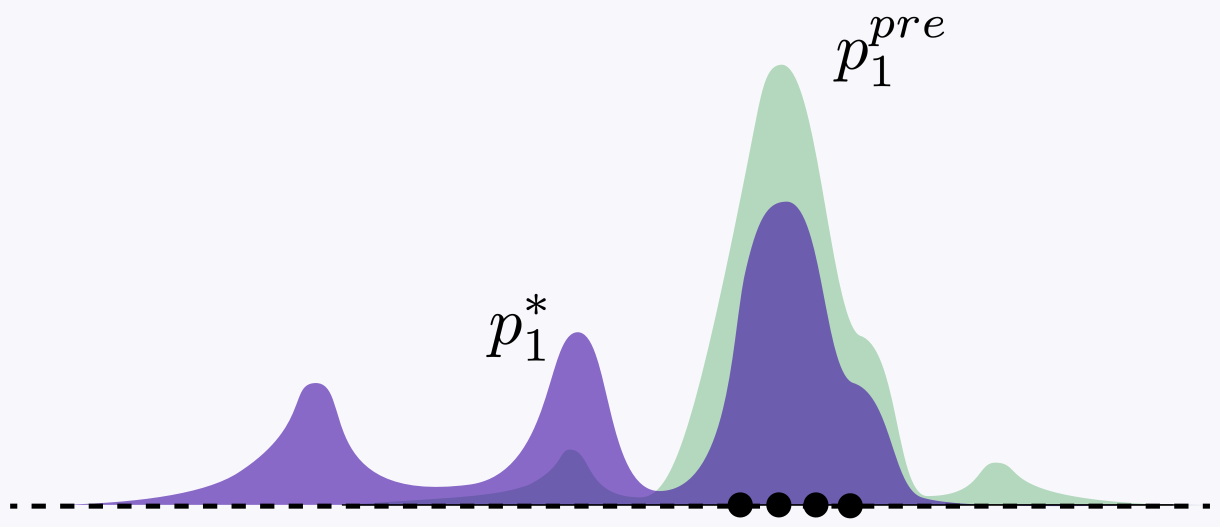

Part III

Out-of-Distribution Flow Modeling

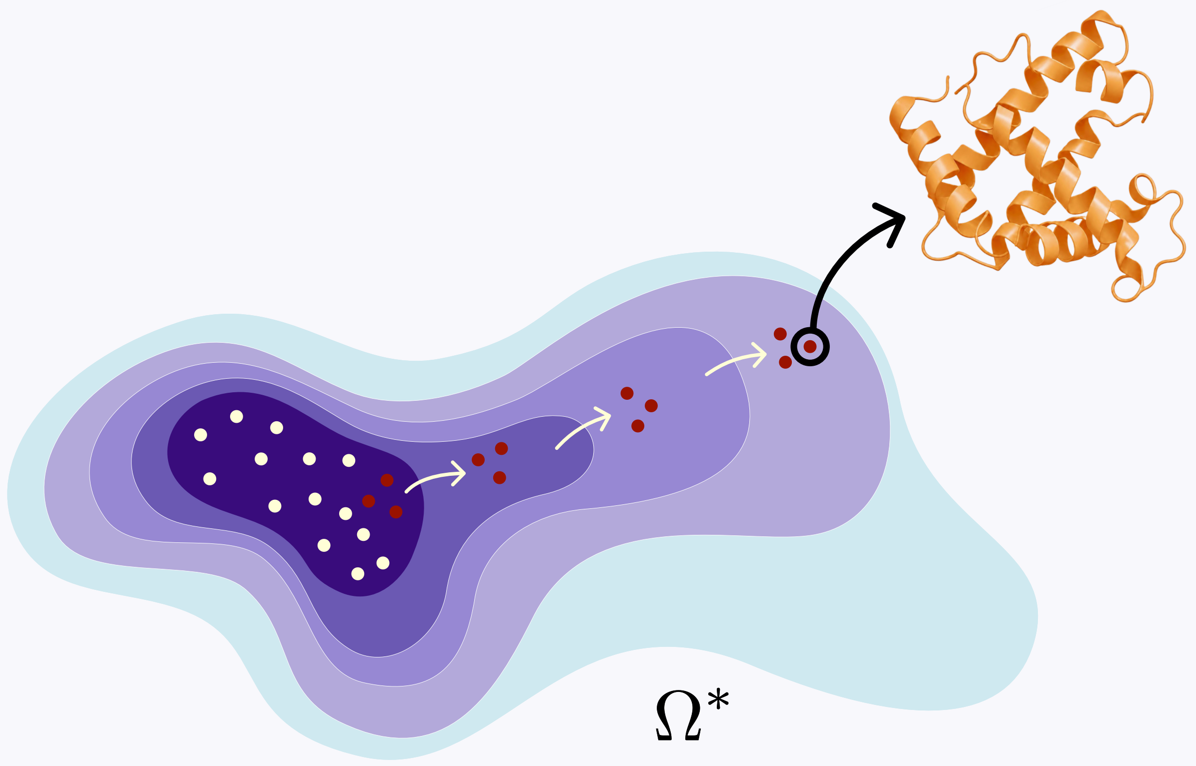

Thanks!

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, Michael M. Zavlanos, Harini Vijayshankar, Dylan Abramson, Basile Wicky, Karl H. Johansson, Kjell Jorner, Florian Dörfler, Pranam Chatterjee, Yisong Yue, Niao He, Andreas Krause.

Funding

Thanks!

Papers + Code

riccardodesanti.com/research

Connect:

X / twitter: @r_de_santi
email: rdesanti@ethz.ch

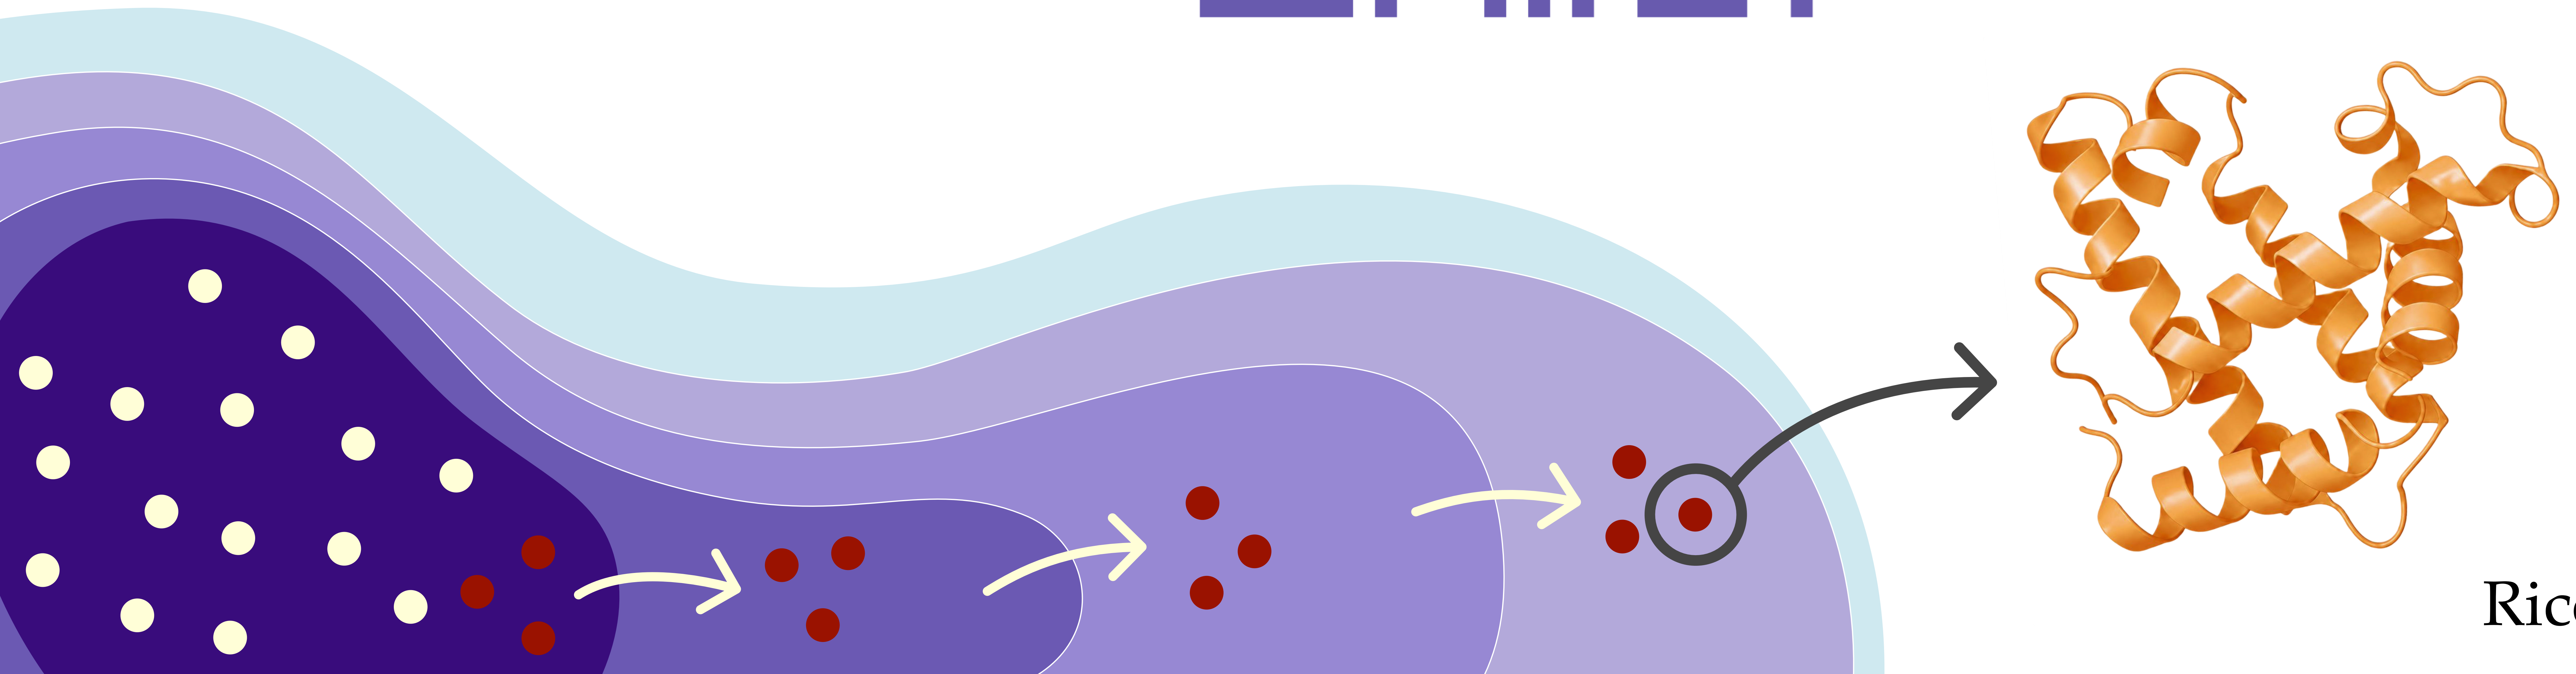

Riccardo De Santi