Epistemic AI Agents for Exploring Chemical Design and Synthesis Spaces

Talk at Technion
Department of Materials Science and Engineering
1.1.2026

Prof. Mario Boley
mboley@is.haifa.ac.il

Department Information Systems
Faculty of Computer and Information Science
University of Haifa

Outline

1. Personas of Agentic AI

2. *Explorer*: mapping polymerisation-induced self-assembly phases

3. *Engineer*: optimising double perovskite hardness

4. *Collector*: discovering novel metal-organic frameworks

5. Problems and directions

AI Agents come in different “Personas”

Agent

at time $t = 1, 2, \ldots$
observes state of world $s_t \in S$
chooses action $a_t \in A$
world transitions to new state $s_{t+1} \sim \mathbb{P}_S(\,\cdot\,; s_t, a_t)$

Reward Maximiser

receives “reward” $r_t \sim \mathbb{P}_R(\,\cdot\,;\, a_t, s_t)$
chooses actions to max. $r_1 + r_2 + \cdots$

Strategist

reinforcement learning

Scholar

state is data
$s_{t+1} = s_t \cup \{(a_t, y_t)\}$

Clerk

transient states
$\boldsymbol{s}_{t+1} \sim \mathbb{P}_S(\,\cdot\,)$

contextual bandits (chat bots)

Gambler

stateless
$s_{t+1} = s_t = \cdots = \emptyset$

multiarmed bandits

Explorer

active learning

Collector

Engineer

blackbox optimisation

Epistemic Agents aiming to Maximise Knowledge

Epistemic Agent: State = Data acquired by Actions

World behaviour captured in conditional distributions

$$y \mid x$$

of

- target variable y given
- input variable $x \in X$

Actions $A = X^k$ allow agent to accrue data $s_0 \subset s_1 \subset s_2 \subset \cdots$

action a_2

x_1	y_1
...	...
x_k	y_k
x_{k+1}	y_{k+1}
...	...
x_{2k}	y_{2k}
...	...
x_{tk}	y_{tk}

state s_t

Agent chooses actions based on response model and uncertainty

$$y \mid x, \theta$$
$$\theta \mid s$$

as usual, to maximise rewards...

Explorer

$r_t = H(\theta \mid s_{t-1}) - H(\theta \mid s_t)$

Collector

$r_t = \#\{(z,y) \in Z \times Y_+ : (z,v,y) \in s_t \setminus s_{t-1}\}$

Engineer

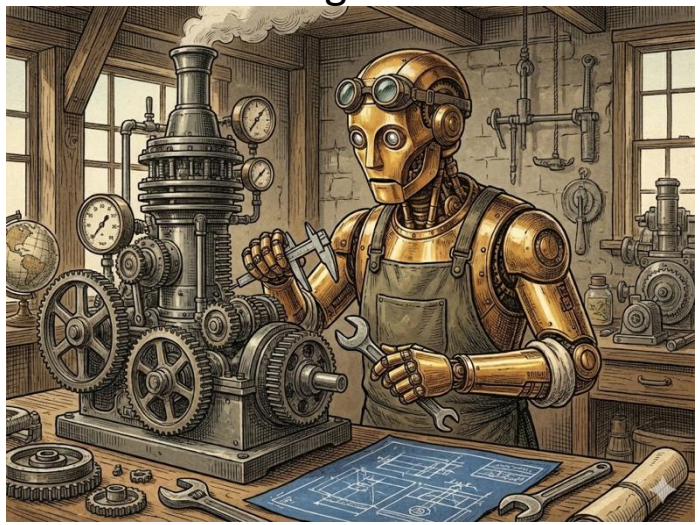

$r_t = \max\{y : (x,y) \in s_t\} - \max\{y : (x,y) \in s_{t-1}\}$

Outline

1. Personas of Agentic AI

2. *Explorer*: mapping polymerisation-induced self-assembly phases

3. *Engineer*: optimising double perovskite hardness

4. *Collector*: discovering novel metal-organic frameworks

5. Problems and directions

Mapping the “Morphological Wonderland” of Polymerisation-Induced Self Assemblies

Yiwen Lu

Dilek Yalcin

Lewis Blackman

MONASH University

CSIRO

Goal: reduce the number of experiments required to obtain accurate phase diagrams of self assemblies of block copolymers

[Y Lu, D Yalcin, P Pigram, L Blackman, and M Boley. "Interpretable machine learning models for phase prediction in polymerization-induced self-assembly." *Journal of Chem. Inf. and Modeling* 63, no. 11 (2023): 3288-3306.]

Polymer-Induced Self-Assemblies

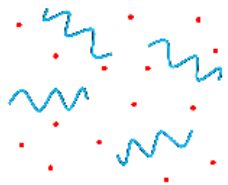

Soluble corona polymer and core monomer

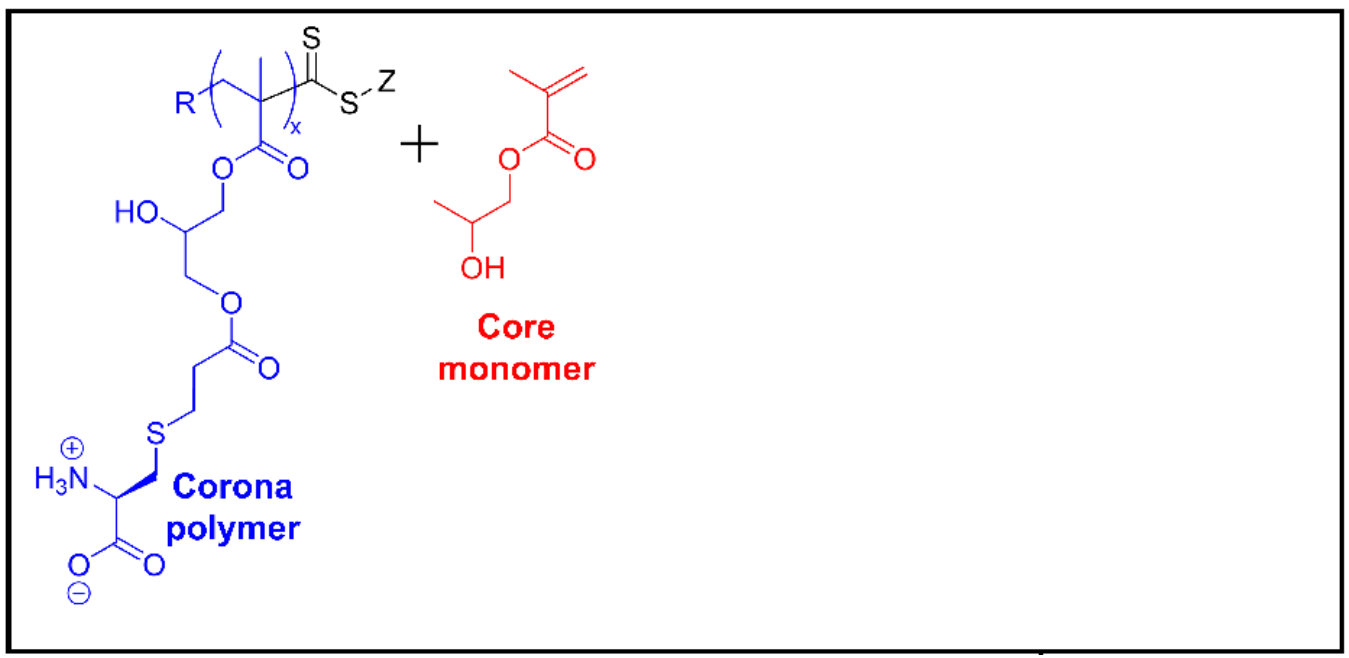

Polymer-Induced Self-Assemblies

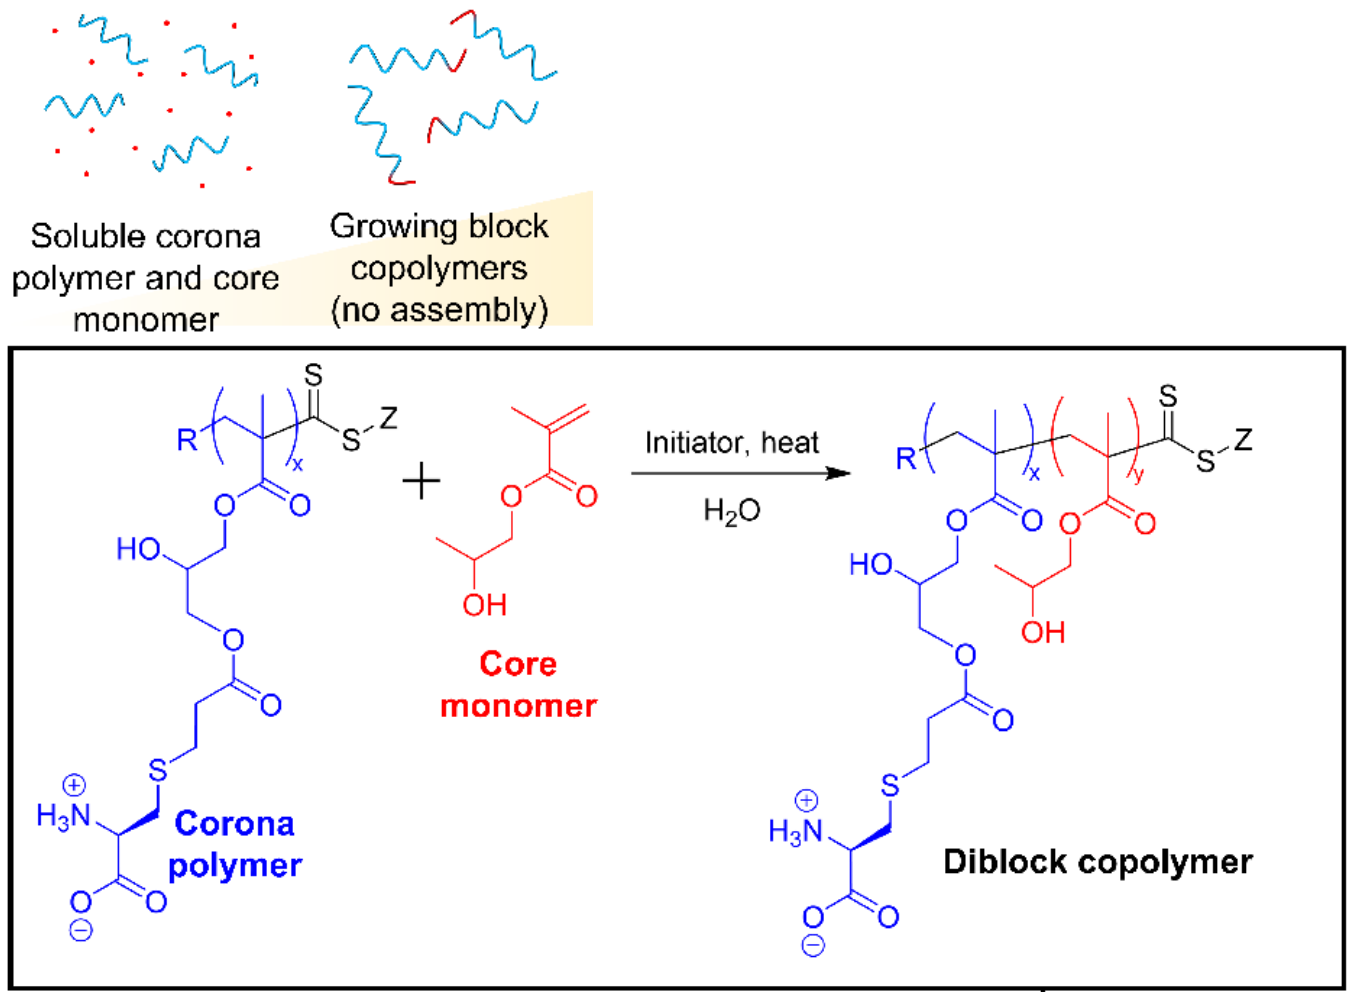

Polymer-Induced Self-Assemblies

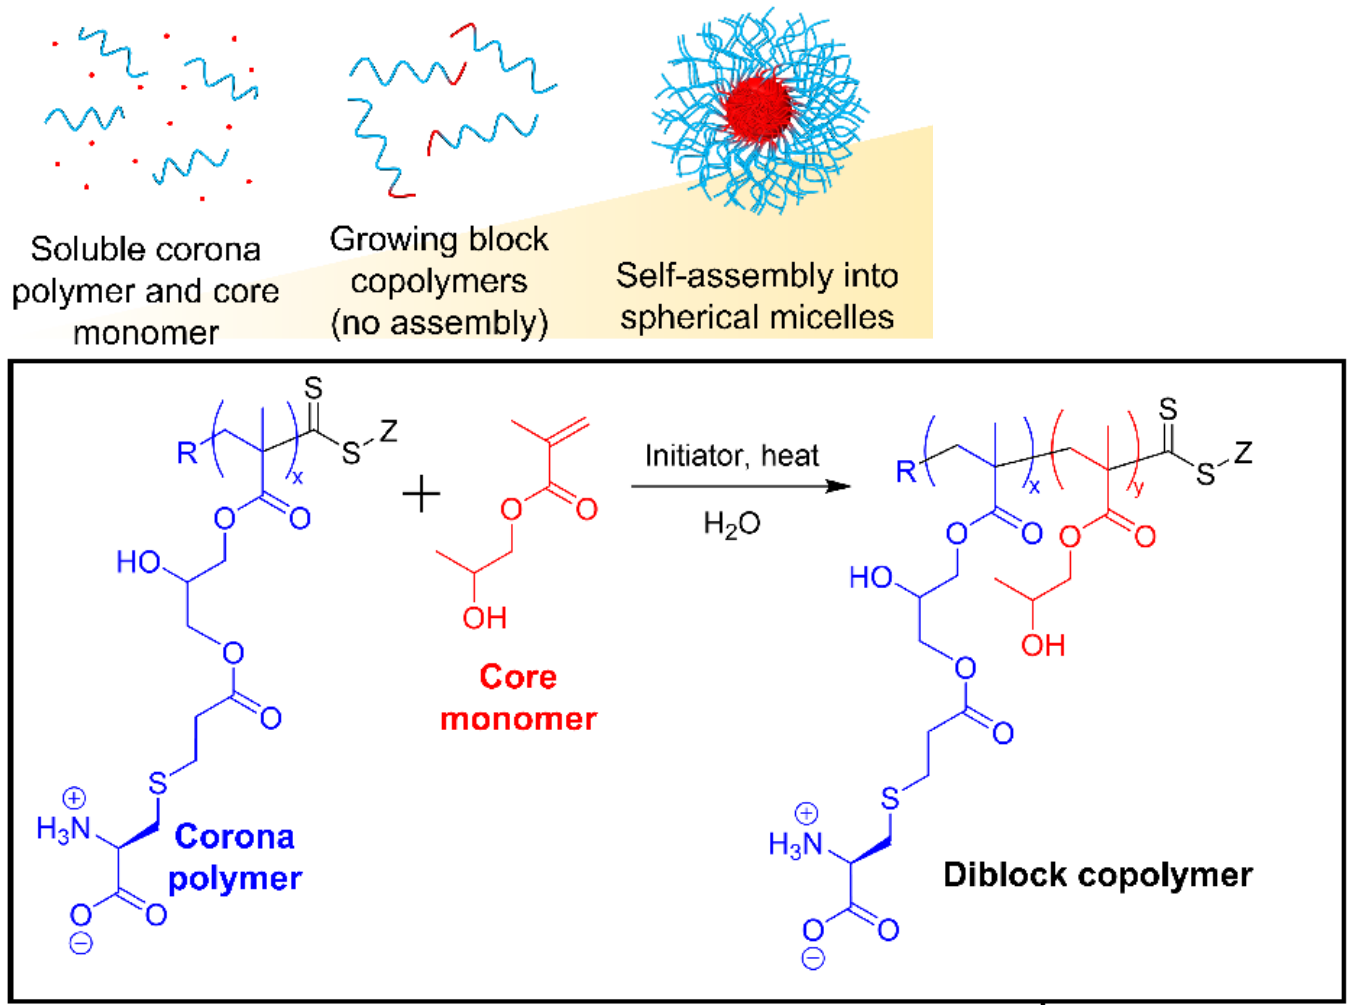

Polymer-Induced Self-Assemblies

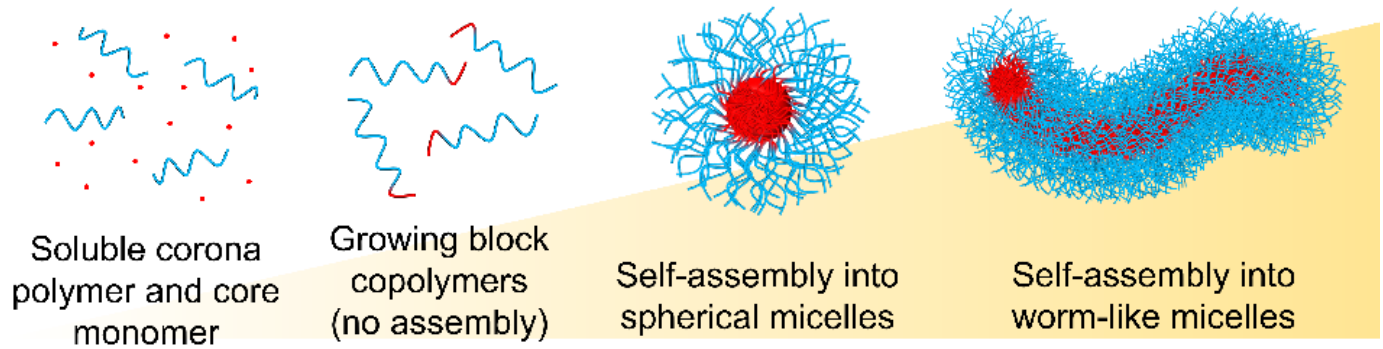

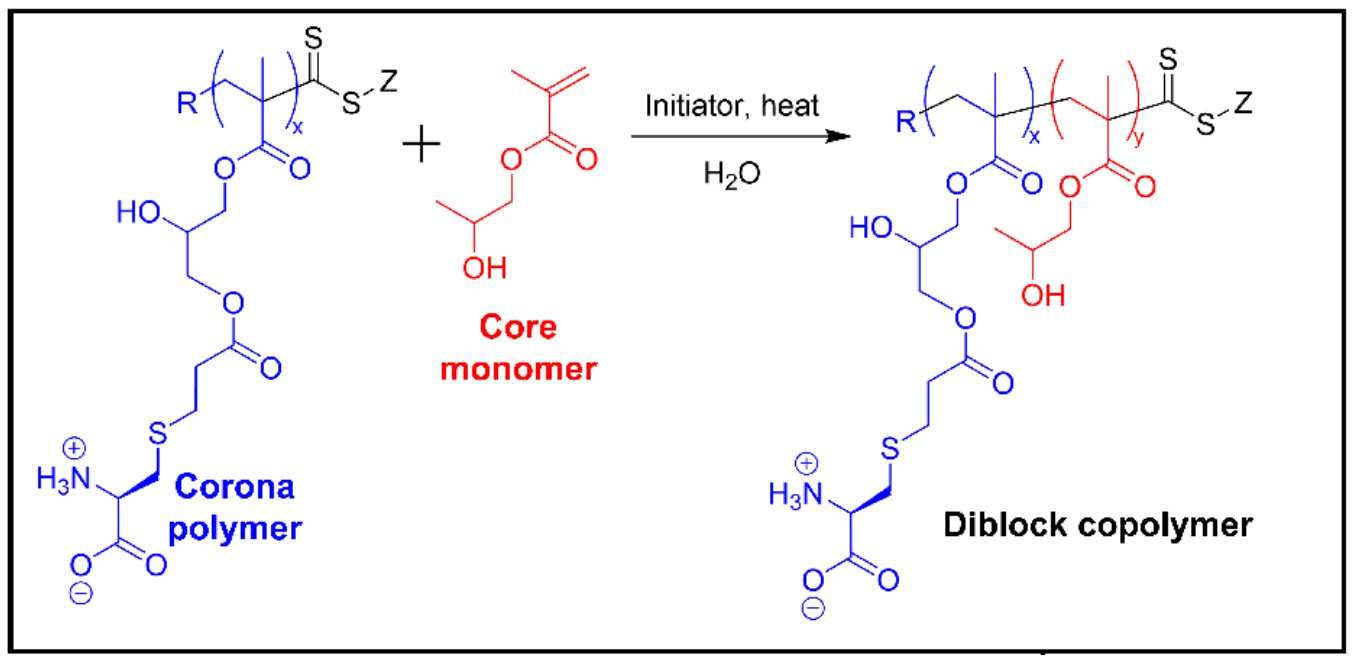

Polymer-Induced Self-Assemblies

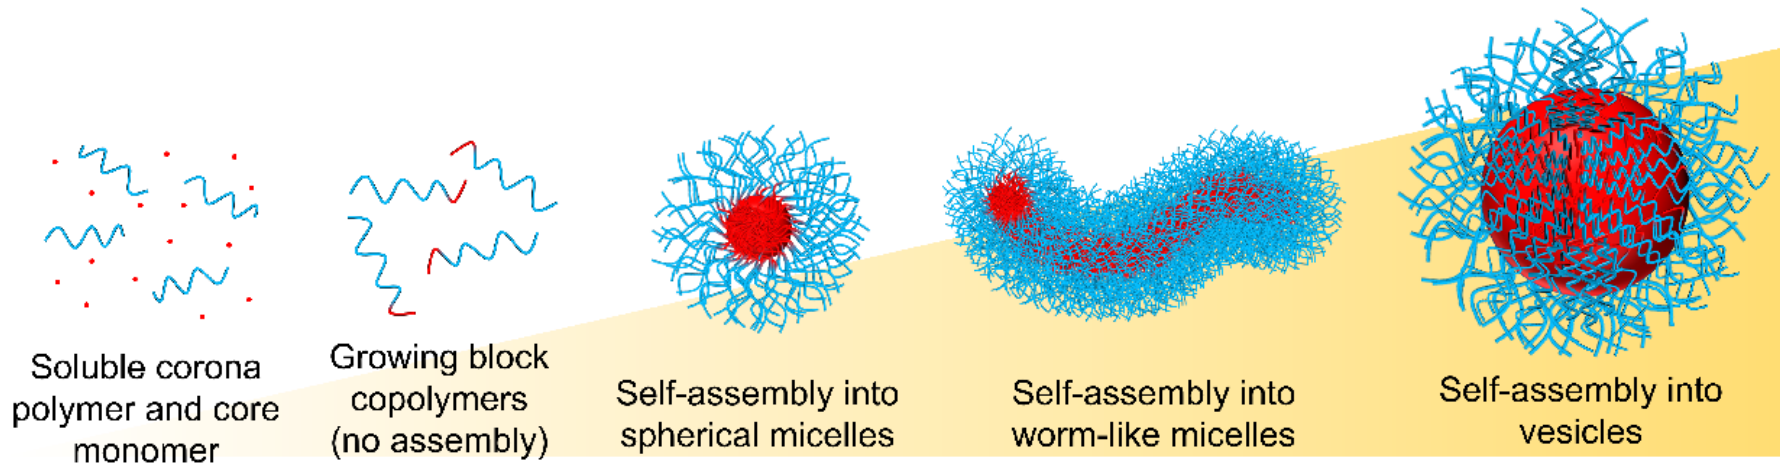

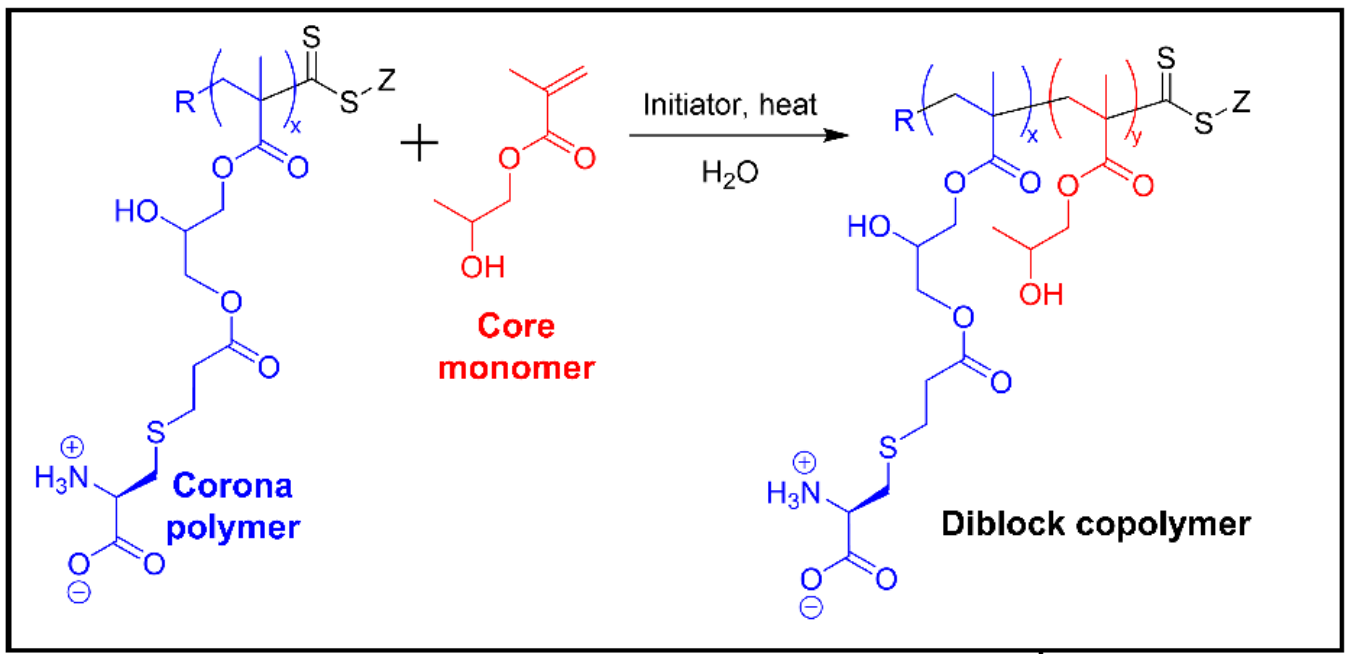

Polymer-Induced Self-Assemblies

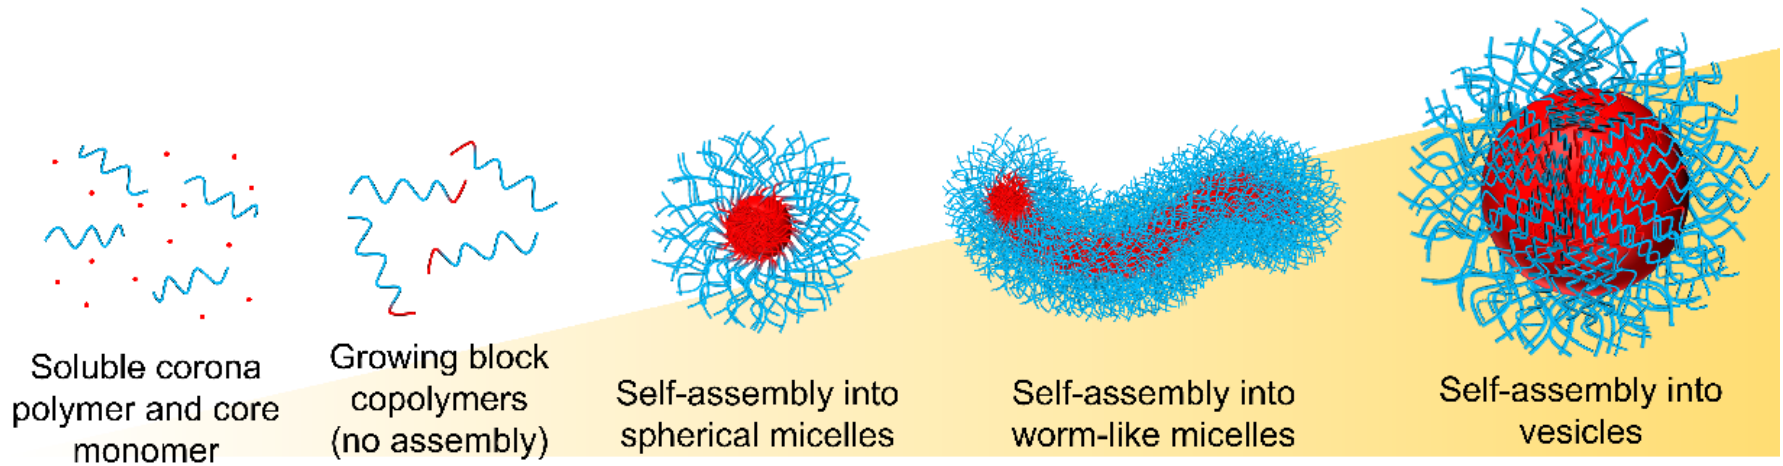

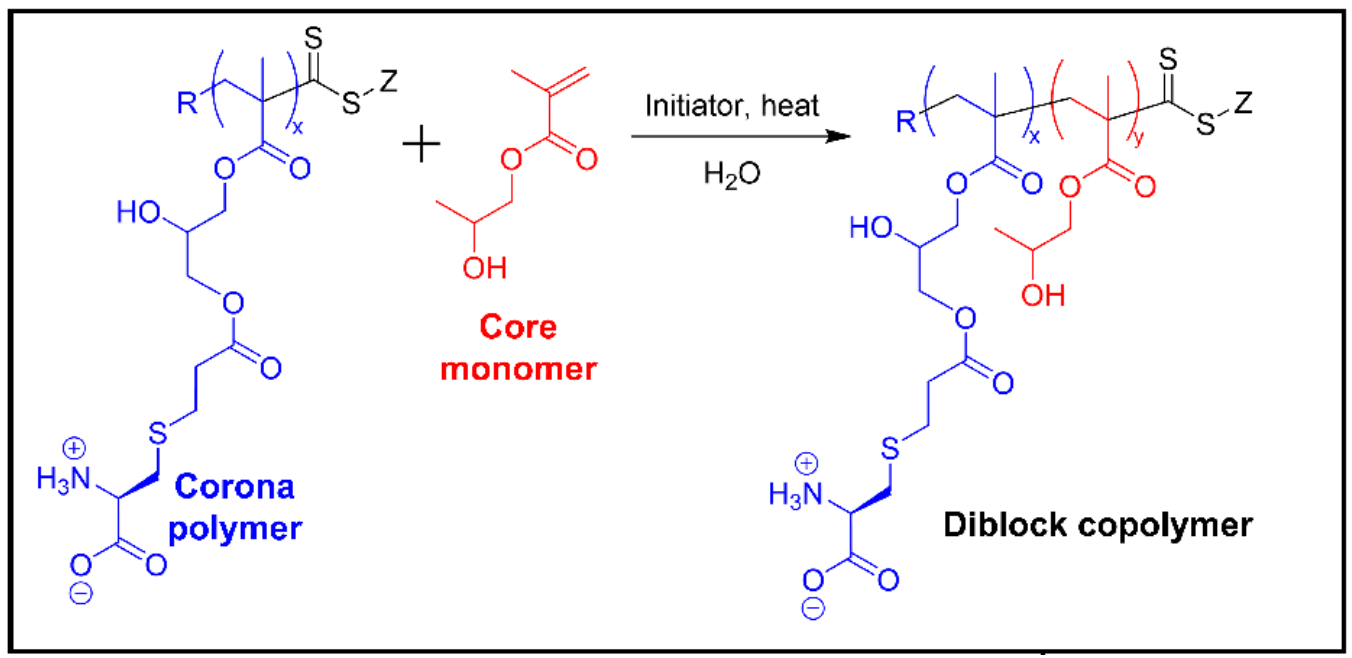

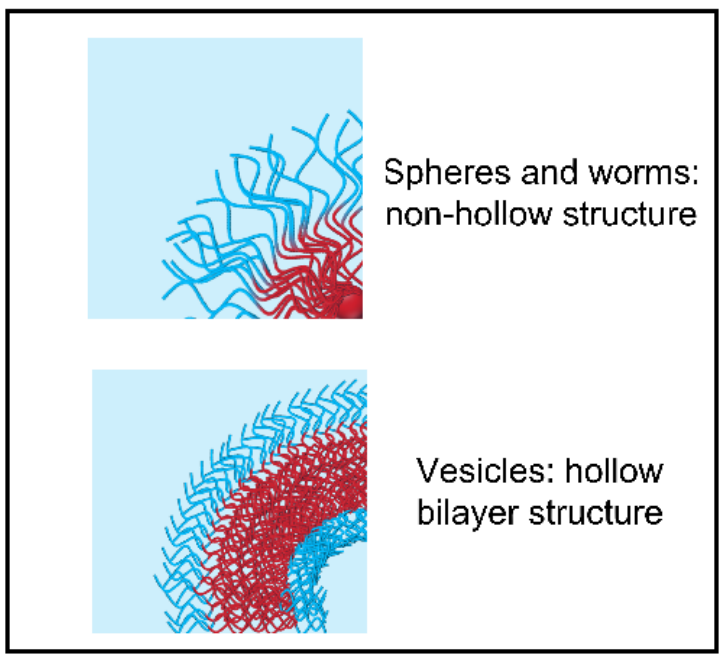

Explorers Minimise Uncertainty about World

Ideas

- Model parameter uncertainty $\theta \mid s$
- Information entropy $H(\theta \mid s)$ captures how much left to learn in state s
- Reward: information gain $H(\theta \mid s_{t-1}) - H(\theta \mid s_t)$
- Data utility: expected information gain

$$u(x) = H(\theta \mid s) - \mathbb{E}\big(H(\theta \mid s \cup \{(x, y)\})\big)$$
$$= H(y \mid x, s) - \mathbb{E}\big(H(y \mid x, \theta)\big)$$

Toy example: 2D Bayesian logistic regression

Model

$$\boldsymbol{\theta} \sim \mathrm{N}_2(0, I)$$
$$y \mid \boldsymbol{\theta} \sim \mathrm{Bern}\big(\sigma(\boldsymbol{x}^T \boldsymbol{\theta})\big)$$

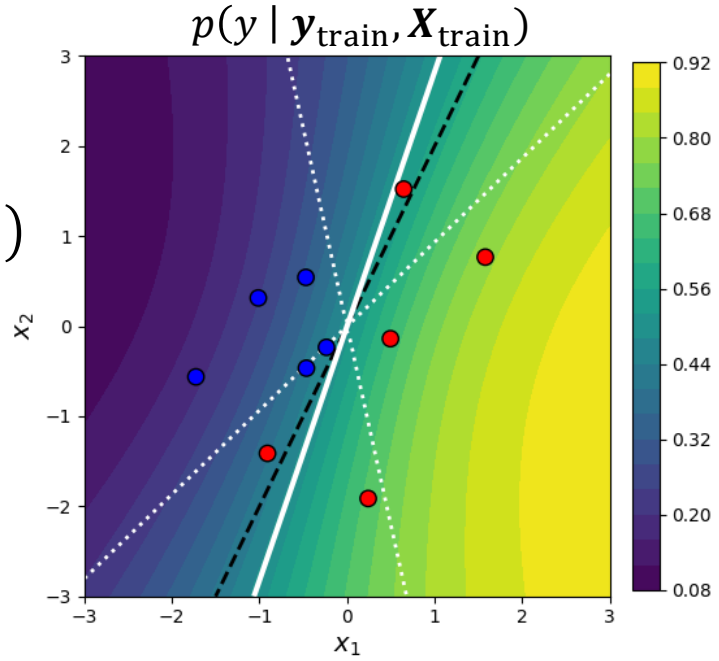

Explorers Minimise Uncertainty about World

Ideas

- Model parameter uncertainty $\theta \mid s$
- Information entropy $H(\theta \mid s)$ captures how much left to learn in state s
- Reward: information gain $H(\theta \mid s_{t-1}) - H(\theta \mid s_t)$
- Data utility: expected information gain

$$u(x) = H(\theta \mid s) - \mathbb{E}\big(H(\theta \mid s \cup \{(x,y)\})\big)$$
$$= H(y \mid x, s) - \mathbb{E}\big(H(y \mid x, \theta)\big)$$

epistemic uncertainty

aleatoric "intrinsic" uncertainty

Toy example: 2D Bayesian logistic regression

Model

$$\boldsymbol{\theta} \sim \mathrm{N}_2(0, I)$$
$$y \mid \boldsymbol{\theta} \sim \mathrm{Bern}\big(\sigma(\boldsymbol{x}^T\boldsymbol{\theta})\big)$$

$p(y \mid \boldsymbol{y}_{\text{train}}, \boldsymbol{X}_{\text{train}})$

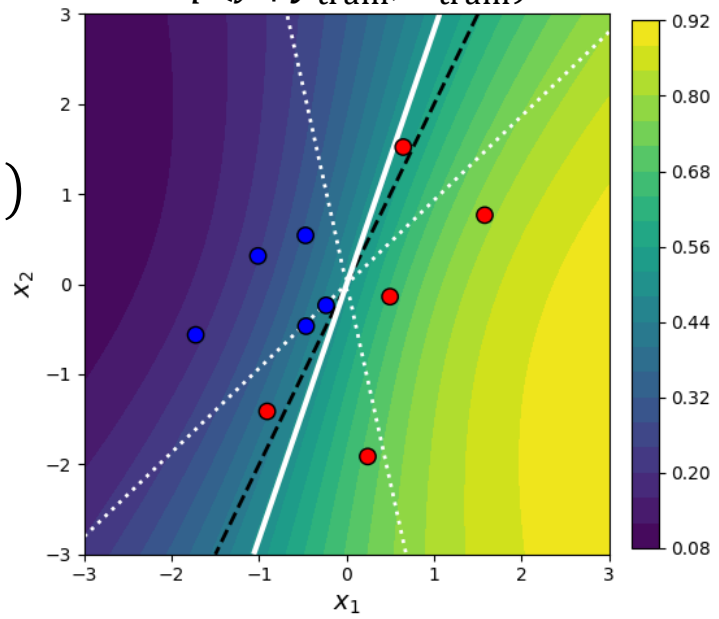

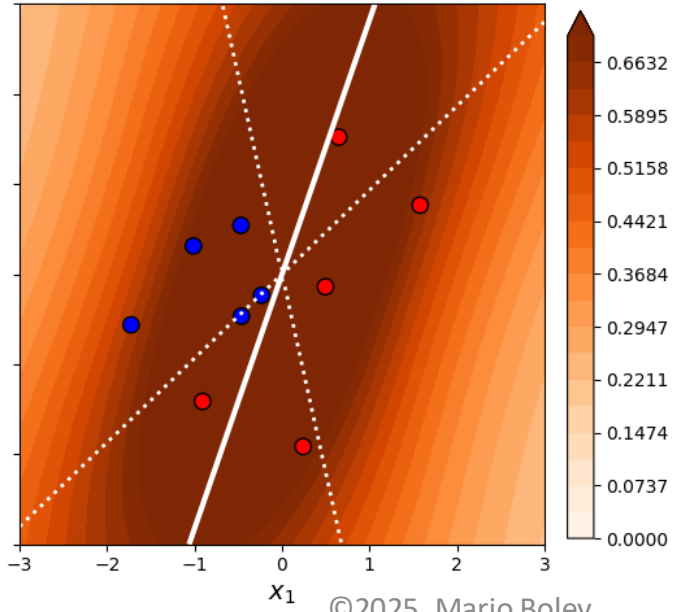

Morphological Outcome Models for PISA

Core and Corona Monomers

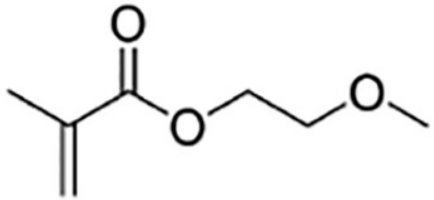

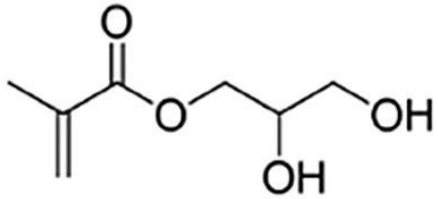

- molar mass
- molecular volume
- polar surface area
- sum of atomic polarizabilities
- log octanol/water partition coefficient (computed)

Conditions

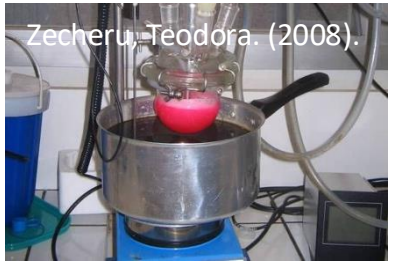

- temperature
- degree of polymerization (core and corona
- concentrations
- solvent pH
- …

Outcomes

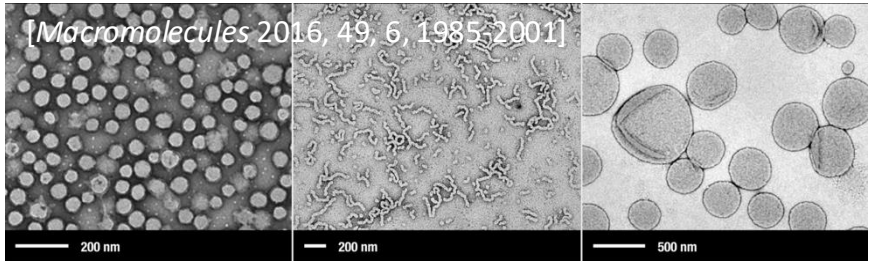

- spheres
- worms
- vesicles
- …

Rapid Phase Diagrams for New Monomer Combinations

Core
HPMA

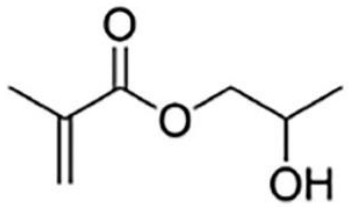

Corona
CysMA

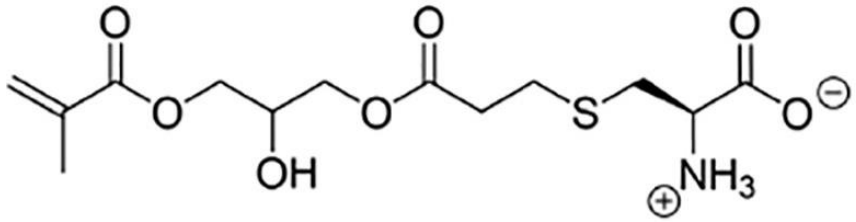

Conditions
DP corona: 31
Temp: 70
pH: 7

$p(y_l \mid x) > 0.5$

initially true only for "sphere" everywhere

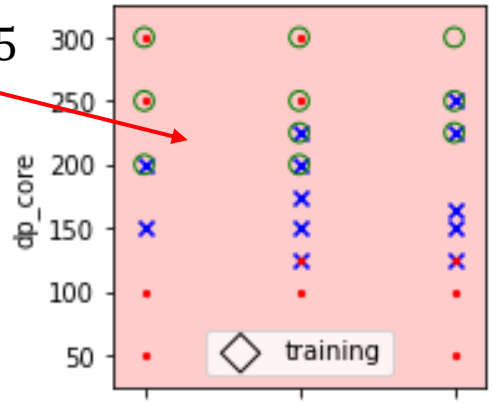

$u(x)$

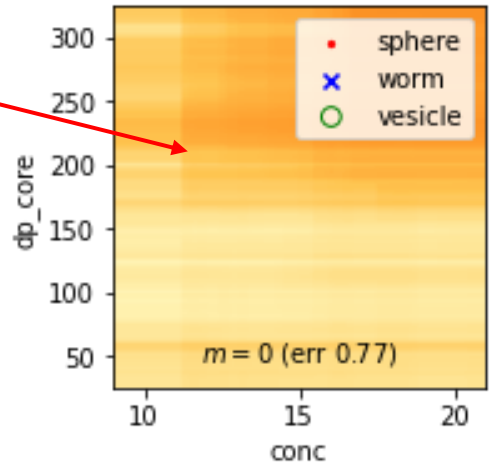

Outline

1. Personas of Agentic AI

2. *Explorer*: mapping polymerisation-induced self-assembly phases

3. *Engineer*: optimising double perovskite hardness

4. *Collector*: discovering novel metal-organic frameworks

5. Problems and directions

Double Perovskite Optimisation as Test Bed for AI-Tools Supporting Rational Materials Design


Simon Teshuva


Phuc “Felix” Luong


Daniel Schmidt


Lucas Foppa


Matthias Scheffler

MONASH University

MAX-PLANCK-GESELLSCHAFT Fritz-Haber-Institut

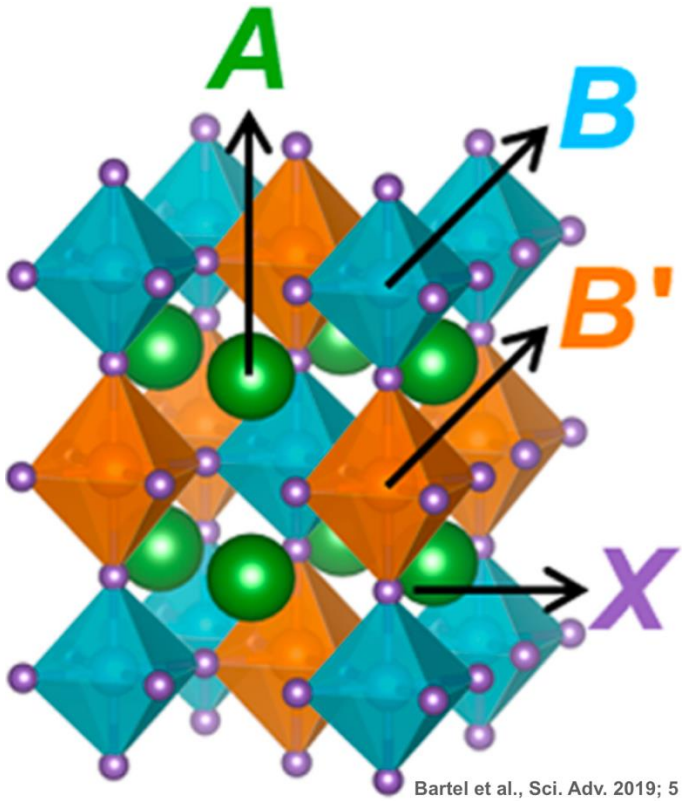

Bartel et al., Sci. Adv. 2019; 5

$A_2BB'X_6$

Goal: development and evaluation of AI-tools for materials property optimization using the example of double perovskites

[M. Boley *et al.*, "From prediction to action: Critical role of performance estimation for machine-learning-driven materials discovery," in S. Bauer *et al.*, *Modelling Simul. Mater. Sci. Eng.* 32, 063301 (2024)]

Engineers Search for Optimal Design

Components

- design space $X = \{x_1, x_2, \dots\}$
- performance $y(x) \in \mathbb{R}$

Reward: improvement

$$r_t = y_t^* - y_{t-1}^*$$
$$y_t^* = \max\{y : (x, y) \in s_t\}$$

Acquisition: expected improvement

$$u(x; s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$

for normal $y \mid x \sim \mathrm{N}(\mu(x), \sigma^2(x))$

$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

"exploitation" "exploration"

Example: double Perovskite oxides (cubic)

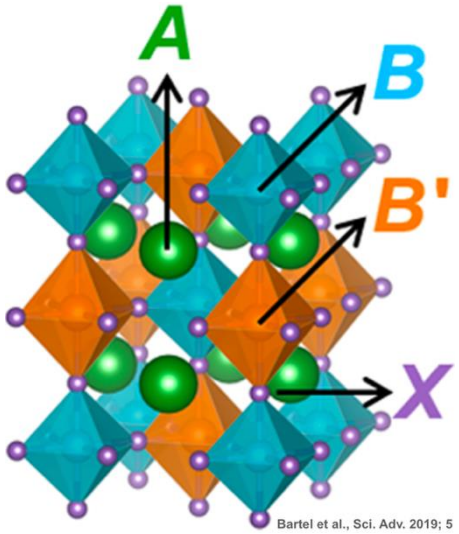

Cations $X_{\mathrm{A}} = \{\mathrm{Li}, \mathrm{Be}, \mathrm{Na}, \mathrm{Mg}, \dots\}$

$X_{\mathrm{B}} = \{\mathrm{Li}, \mathrm{Be}, \mathrm{B}, \mathrm{N}, \dots\}$

$X_{\mathrm{B}'} = \{\mathrm{Be}, \mathrm{B}, \mathrm{C}, \mathrm{N} \dots\}$

Anion $X_{\mathrm{X}} = \{\mathrm{O}\}$

$$X = \{(a, b, b', x) \in X_A \times X_B \times X_{B'} \times X_X : b \le b', \dots\}$$

$y(x) = -V\mathrm{d}P/\mathrm{d}V$: DFT computed bulk modulus

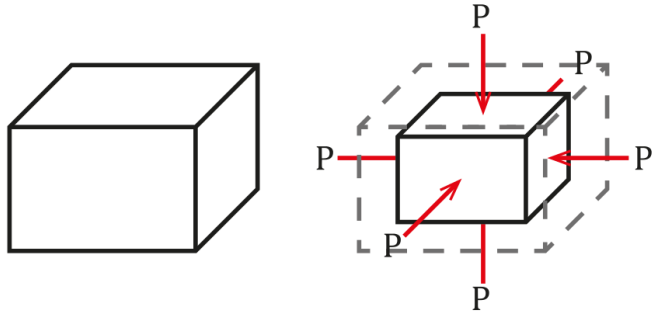

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

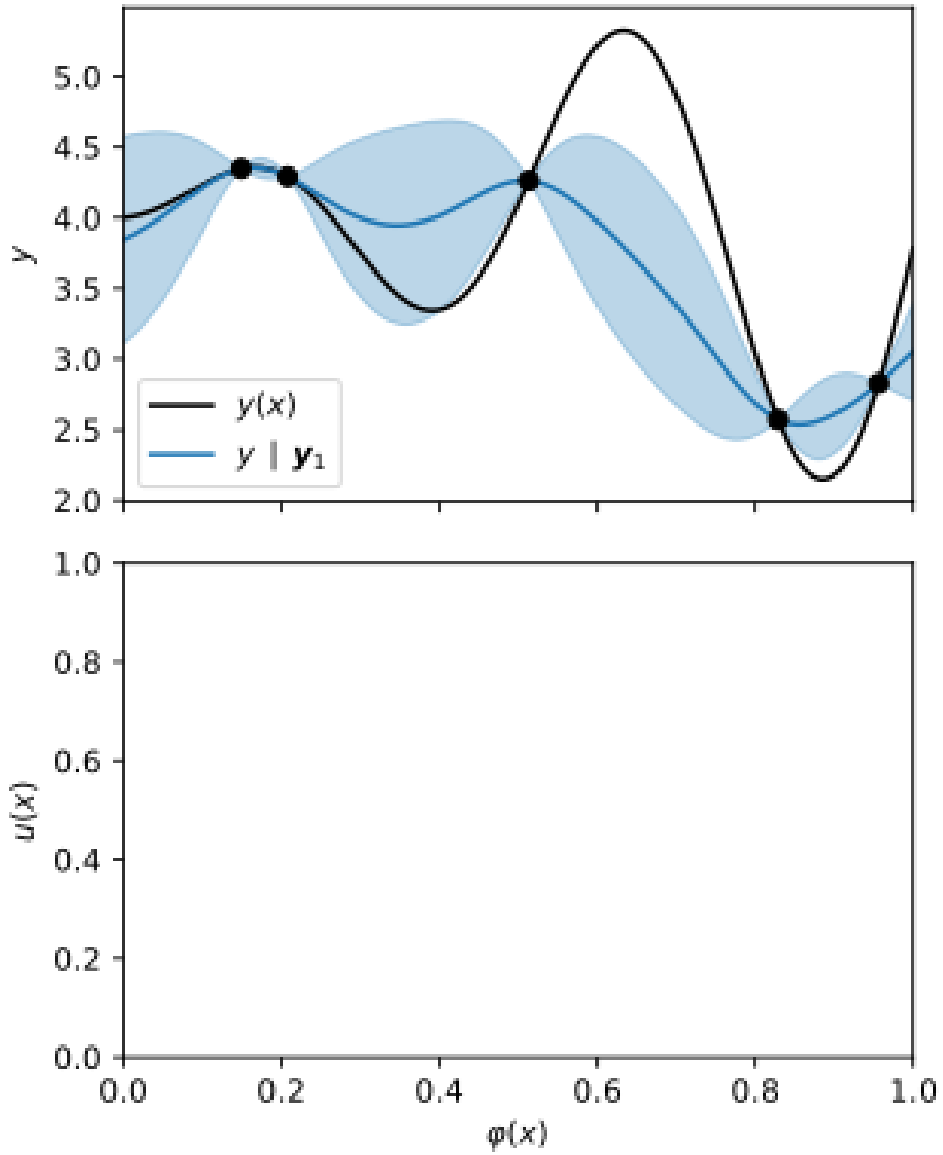

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

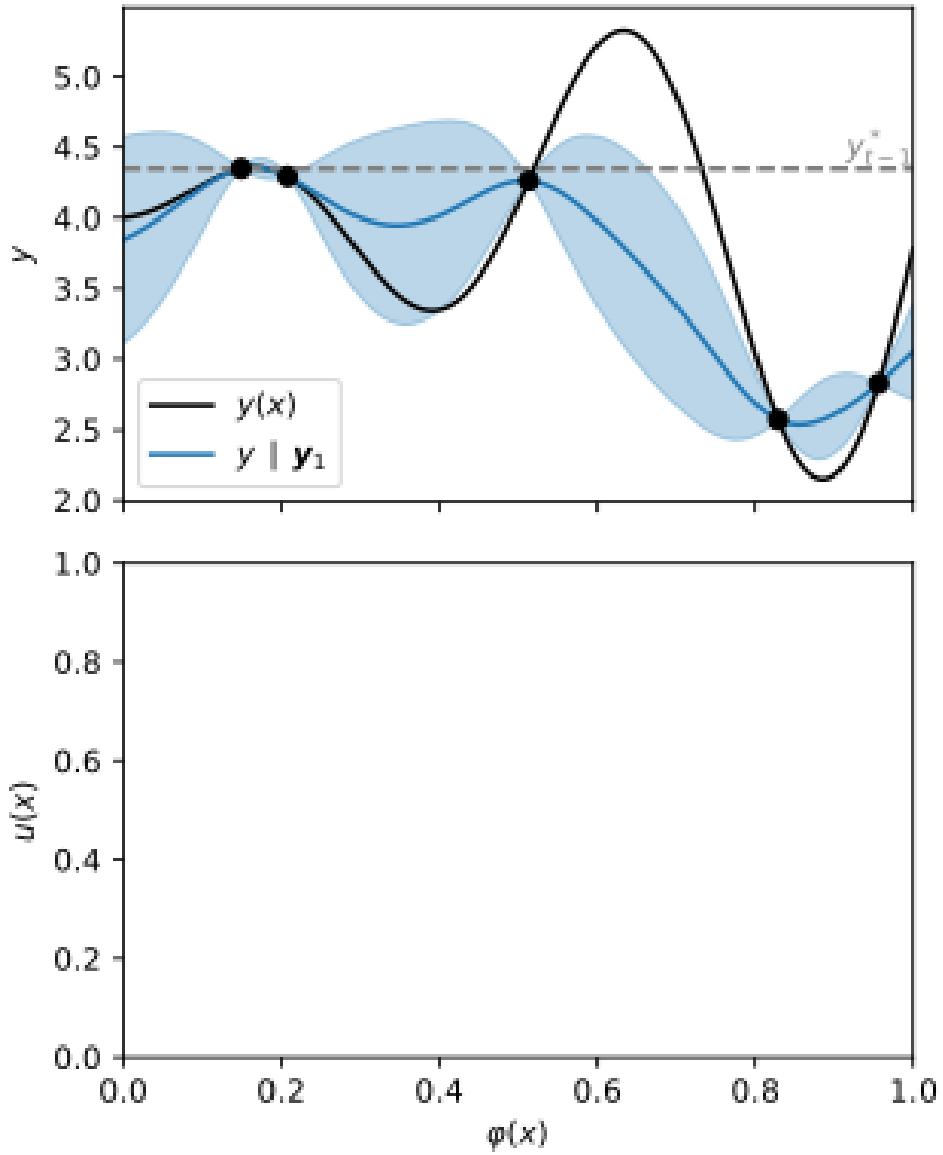

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

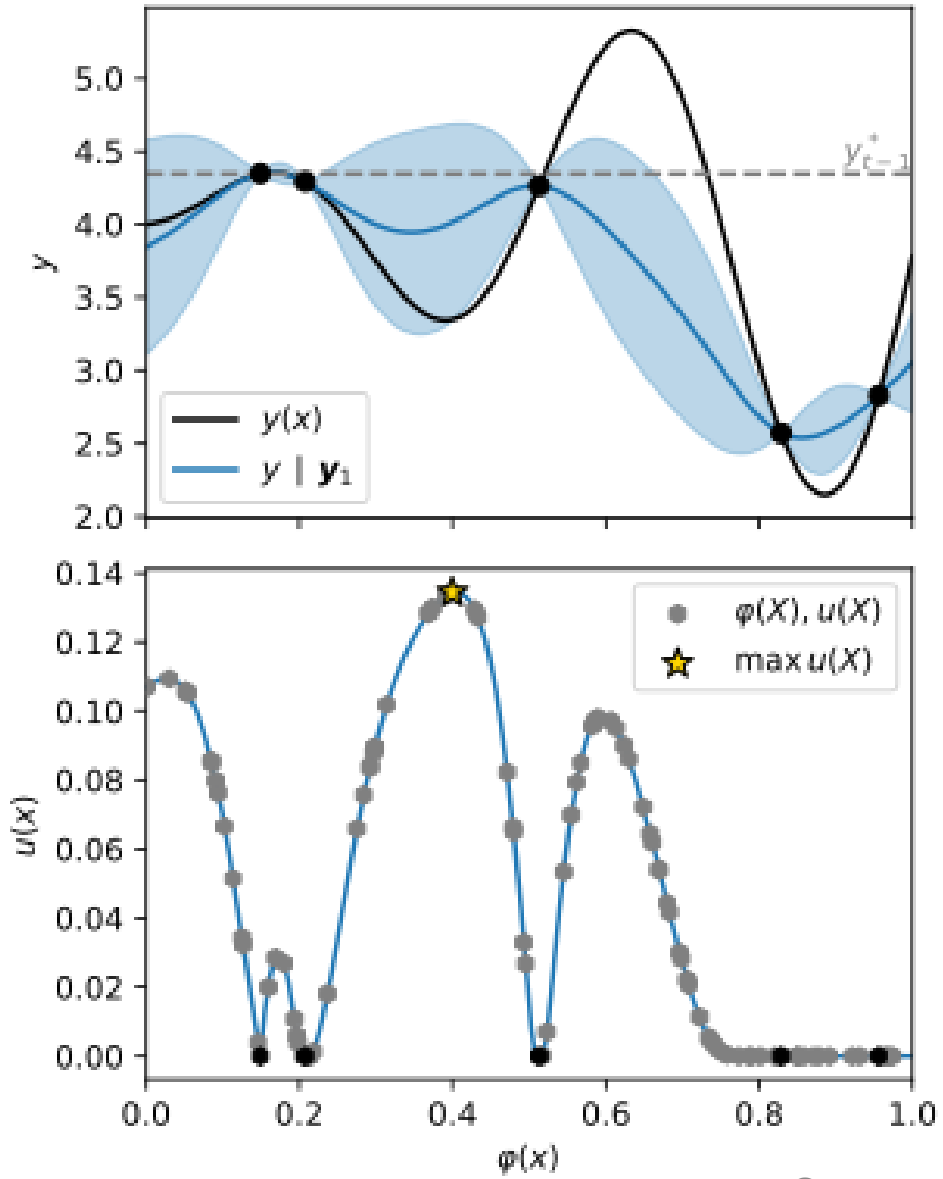

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

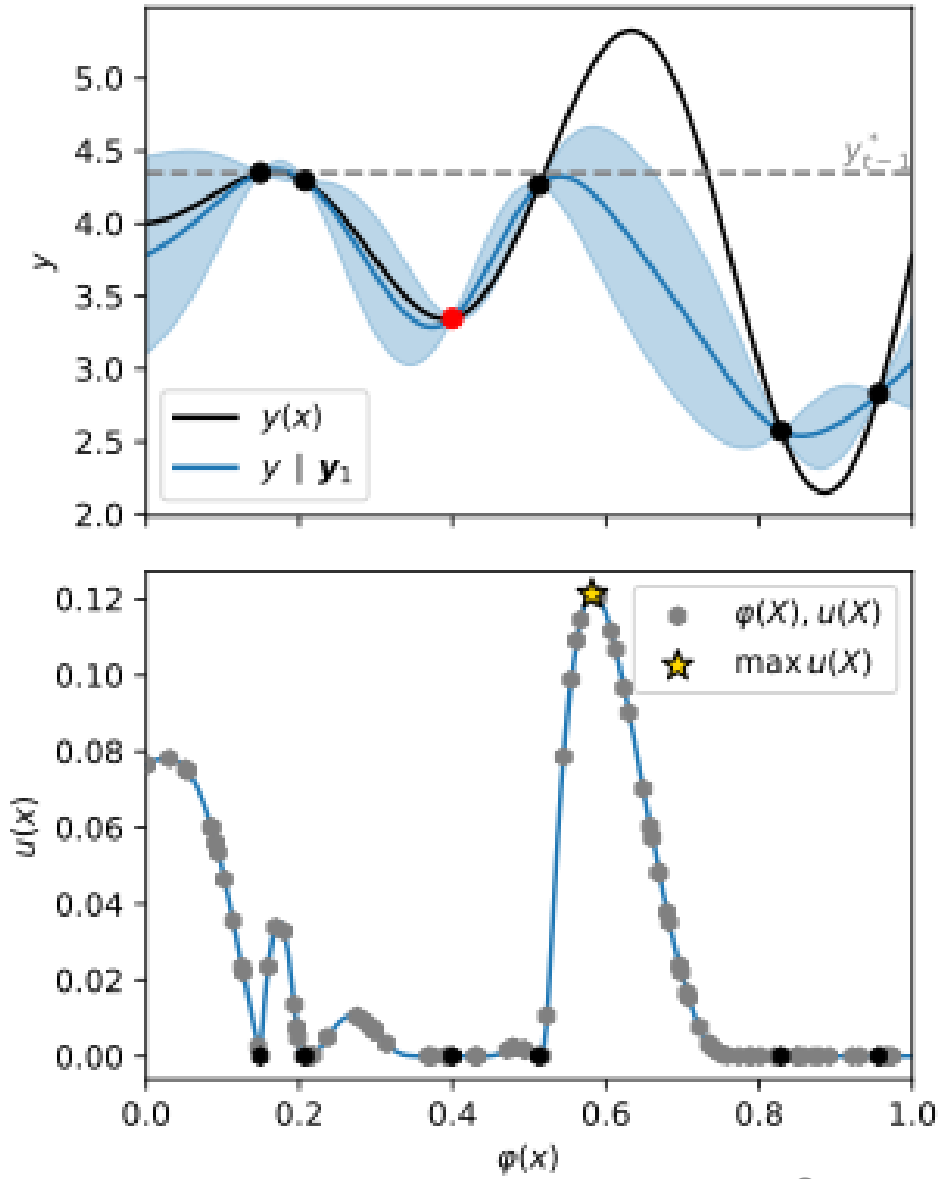

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

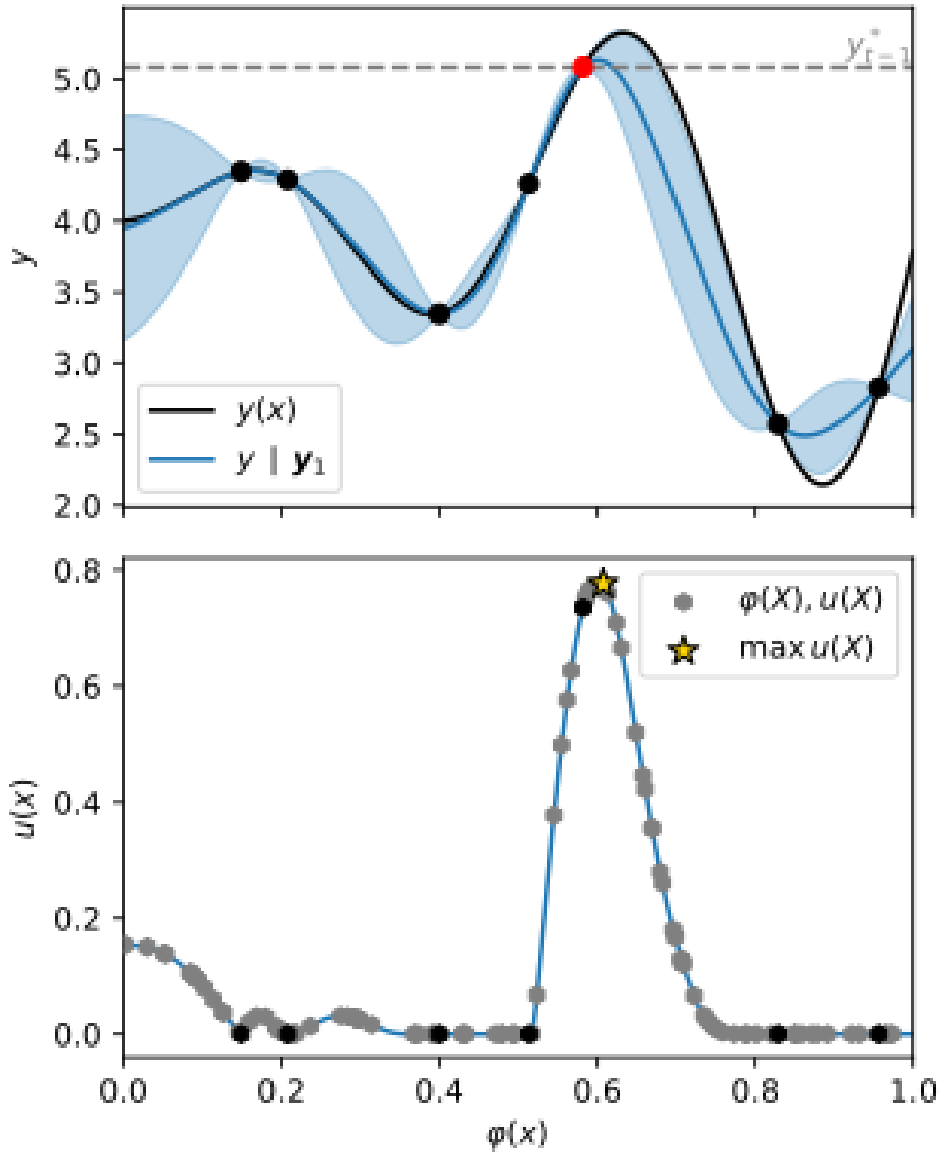

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

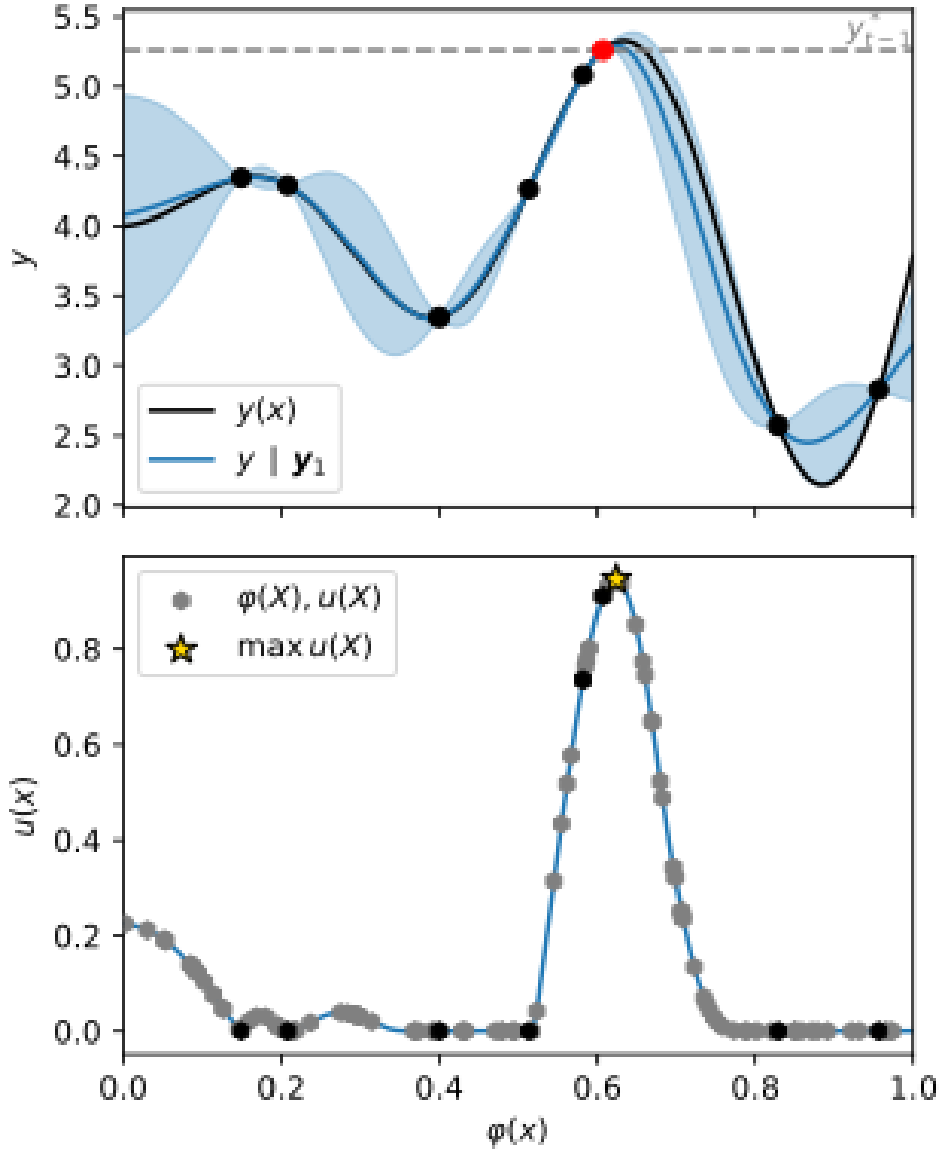

Gaussian Models yield Tractable Reward Estimation

Acquisition: expected improvement

$$u(x;s) = \mathbb{E}(\max\{y - y_s^*, 0\})$$
$$= (\mu(x) - y_s^*)\big(1 - \Phi(y_s^*)\big) + \sigma^2(x)\phi(y_s^*)$$

Gaussian process regression

$$\boldsymbol{y} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K} + \sigma^2\boldsymbol{I})$$
$$\boldsymbol{y}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\big(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\big)$$

$$\boldsymbol{\mu}_2 = \boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\boldsymbol{y}_1$$
$$\boldsymbol{\Sigma}_{2,2} = \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\big(\boldsymbol{K}_{1,1} + \sigma^{-2}\boldsymbol{I}\big)^{-1}\boldsymbol{K}_{1,2}$$

Random forest

$$f \mid \boldsymbol{y}_1, \boldsymbol{X}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$
$$y \mid \boldsymbol{y}_1 \sim \mathrm{N}\Big(\mathbb{E}\big(f(\boldsymbol{x})\big), \mathbb{V}\big(f(\boldsymbol{x})\big)\Big)$$

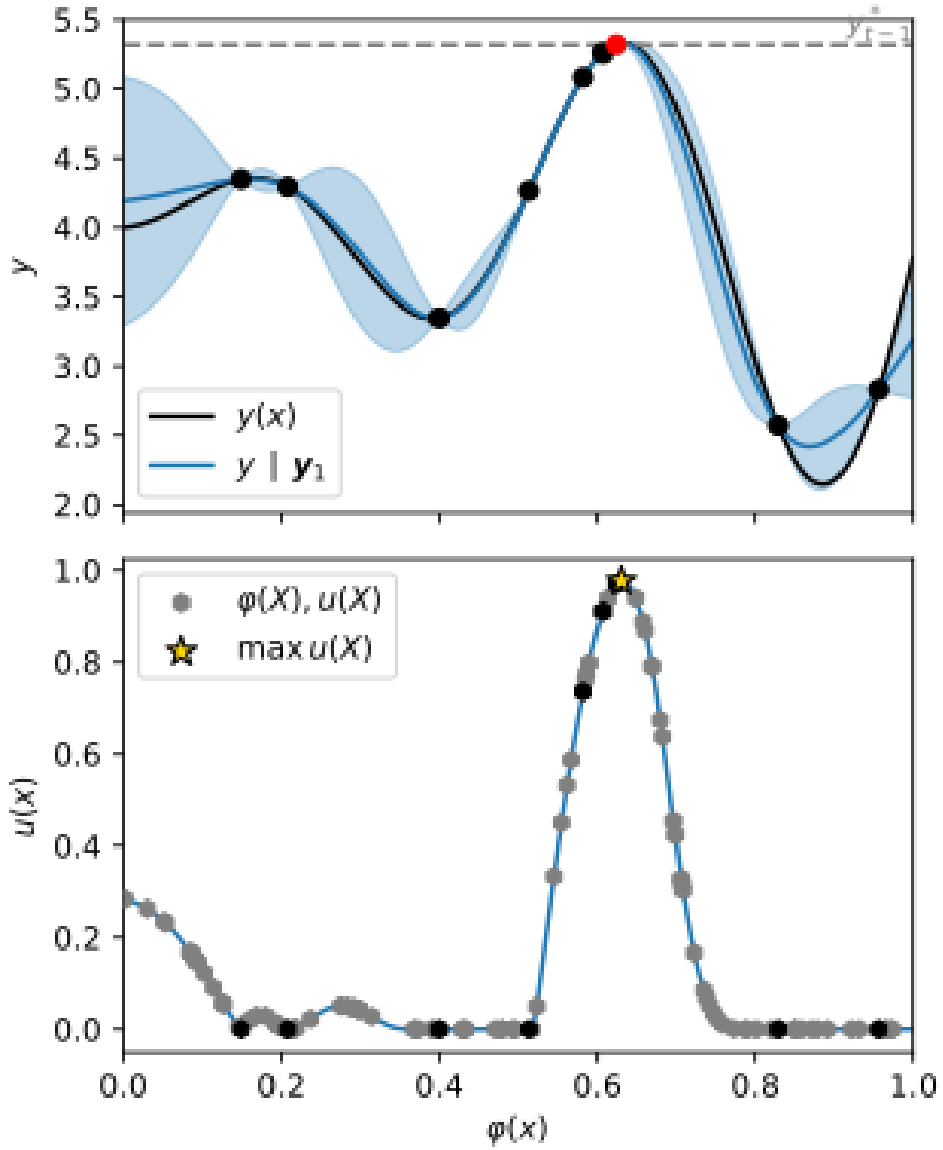

Strategy performs both: Exploration and Exploitation

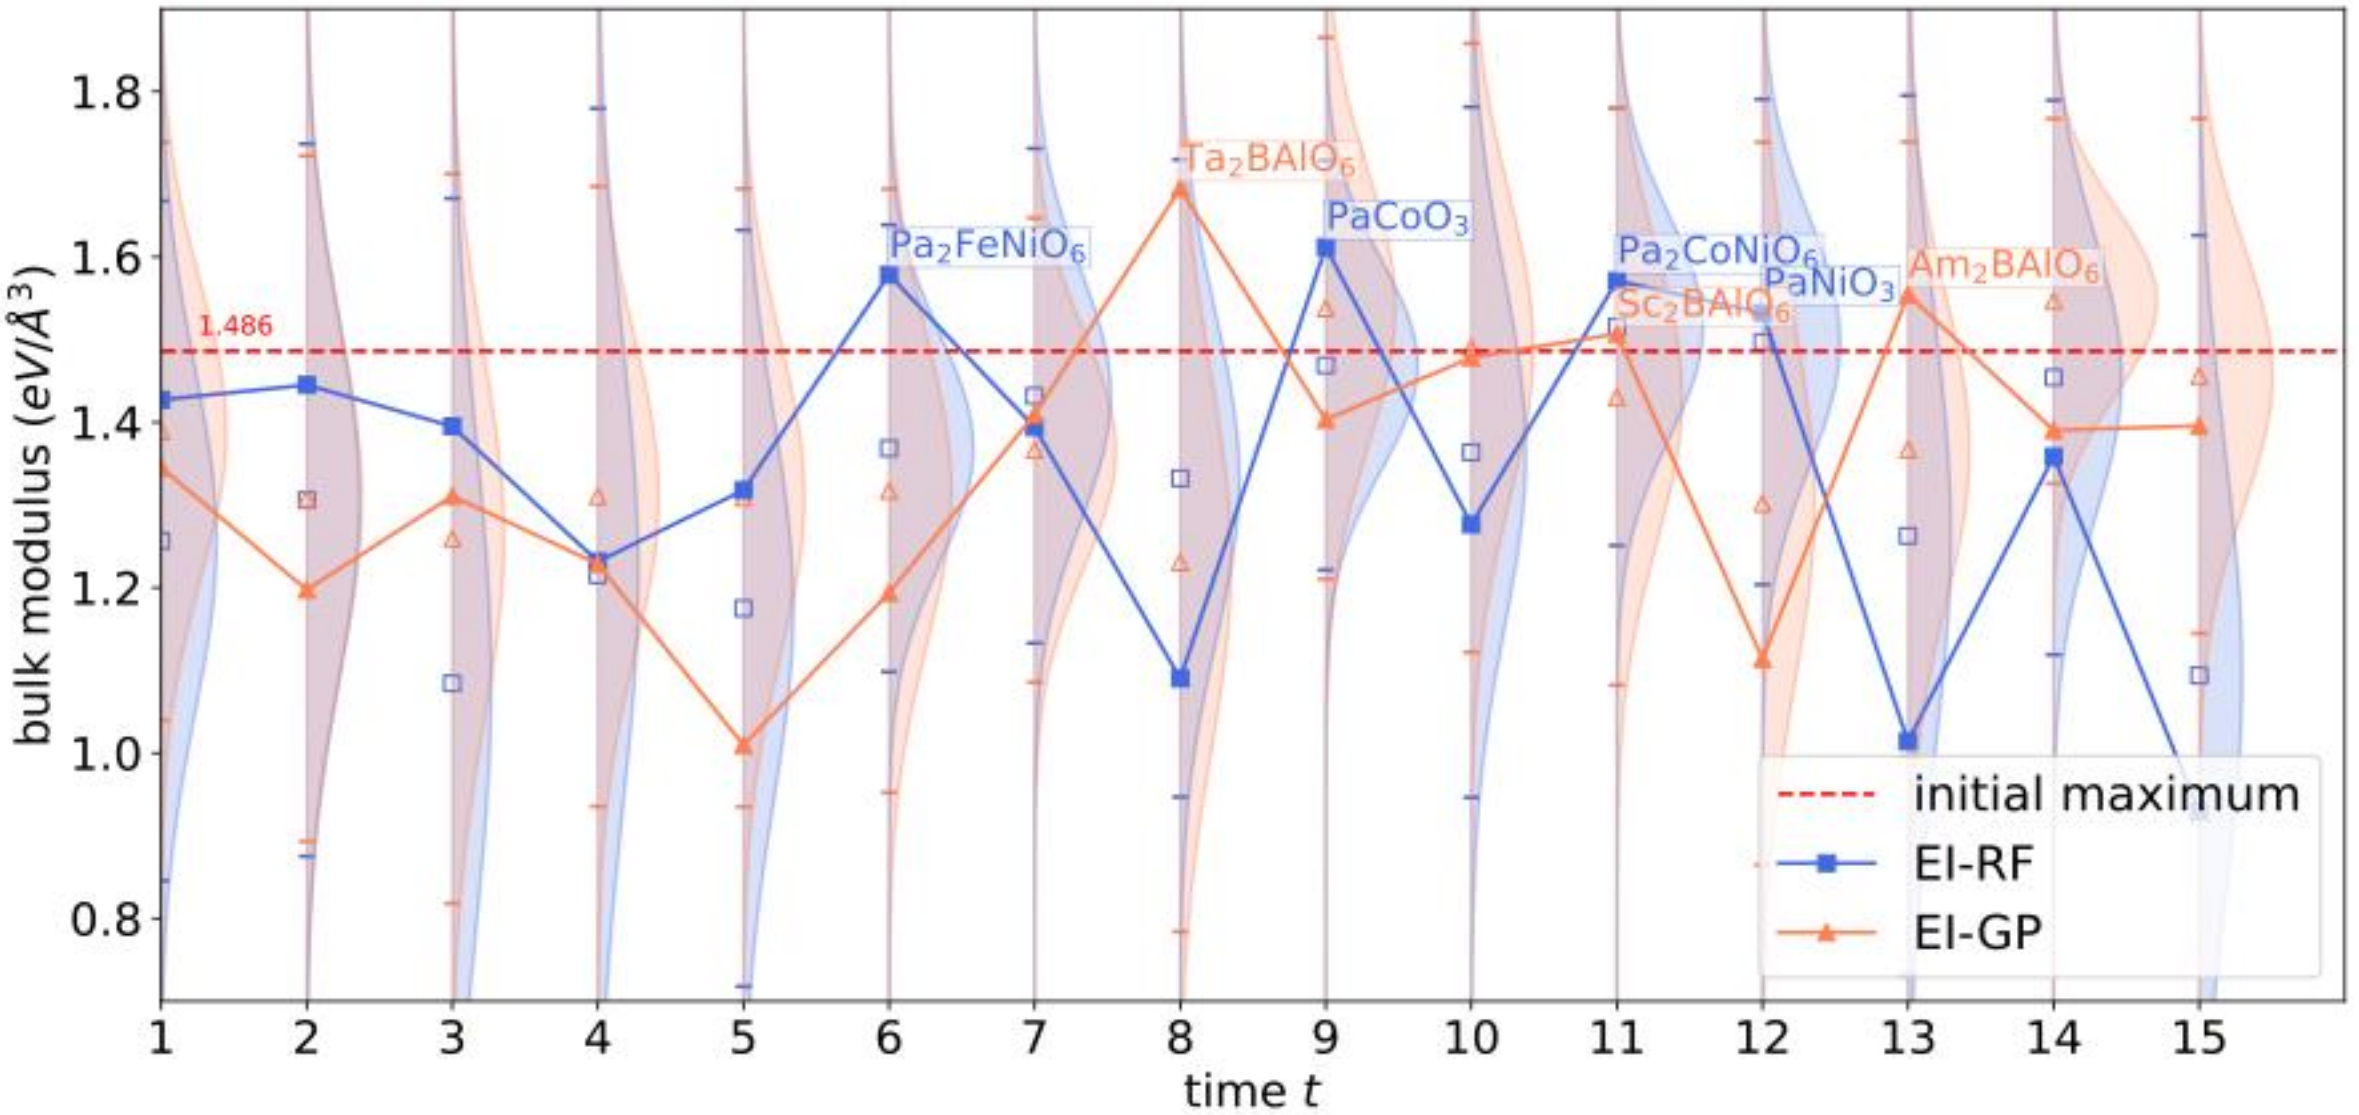

Strategy performs both: Exploration and Exploitation

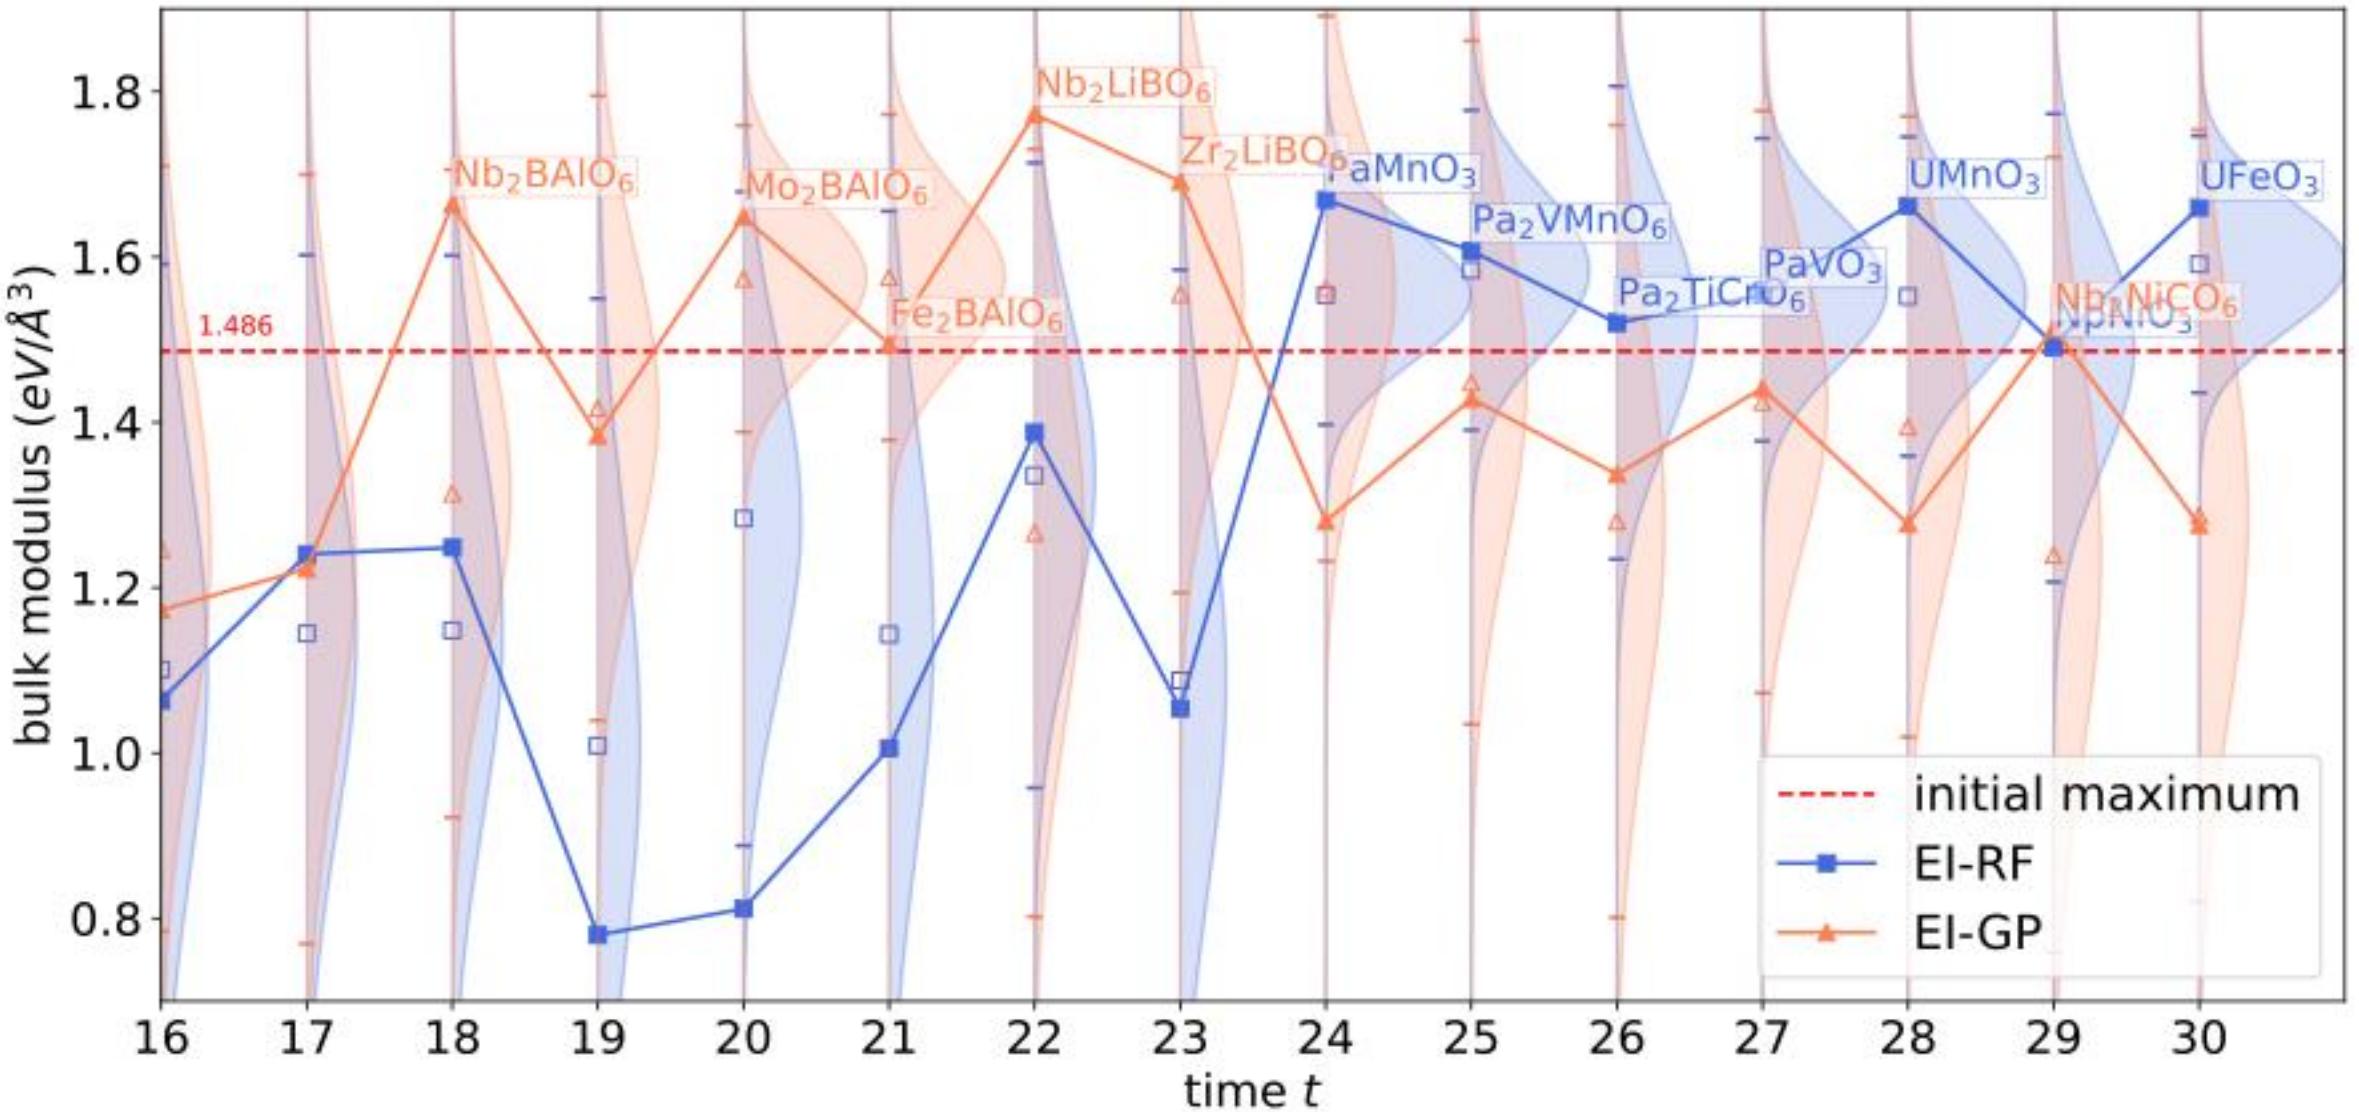

Strategy performs both: Exploration and Exploitation

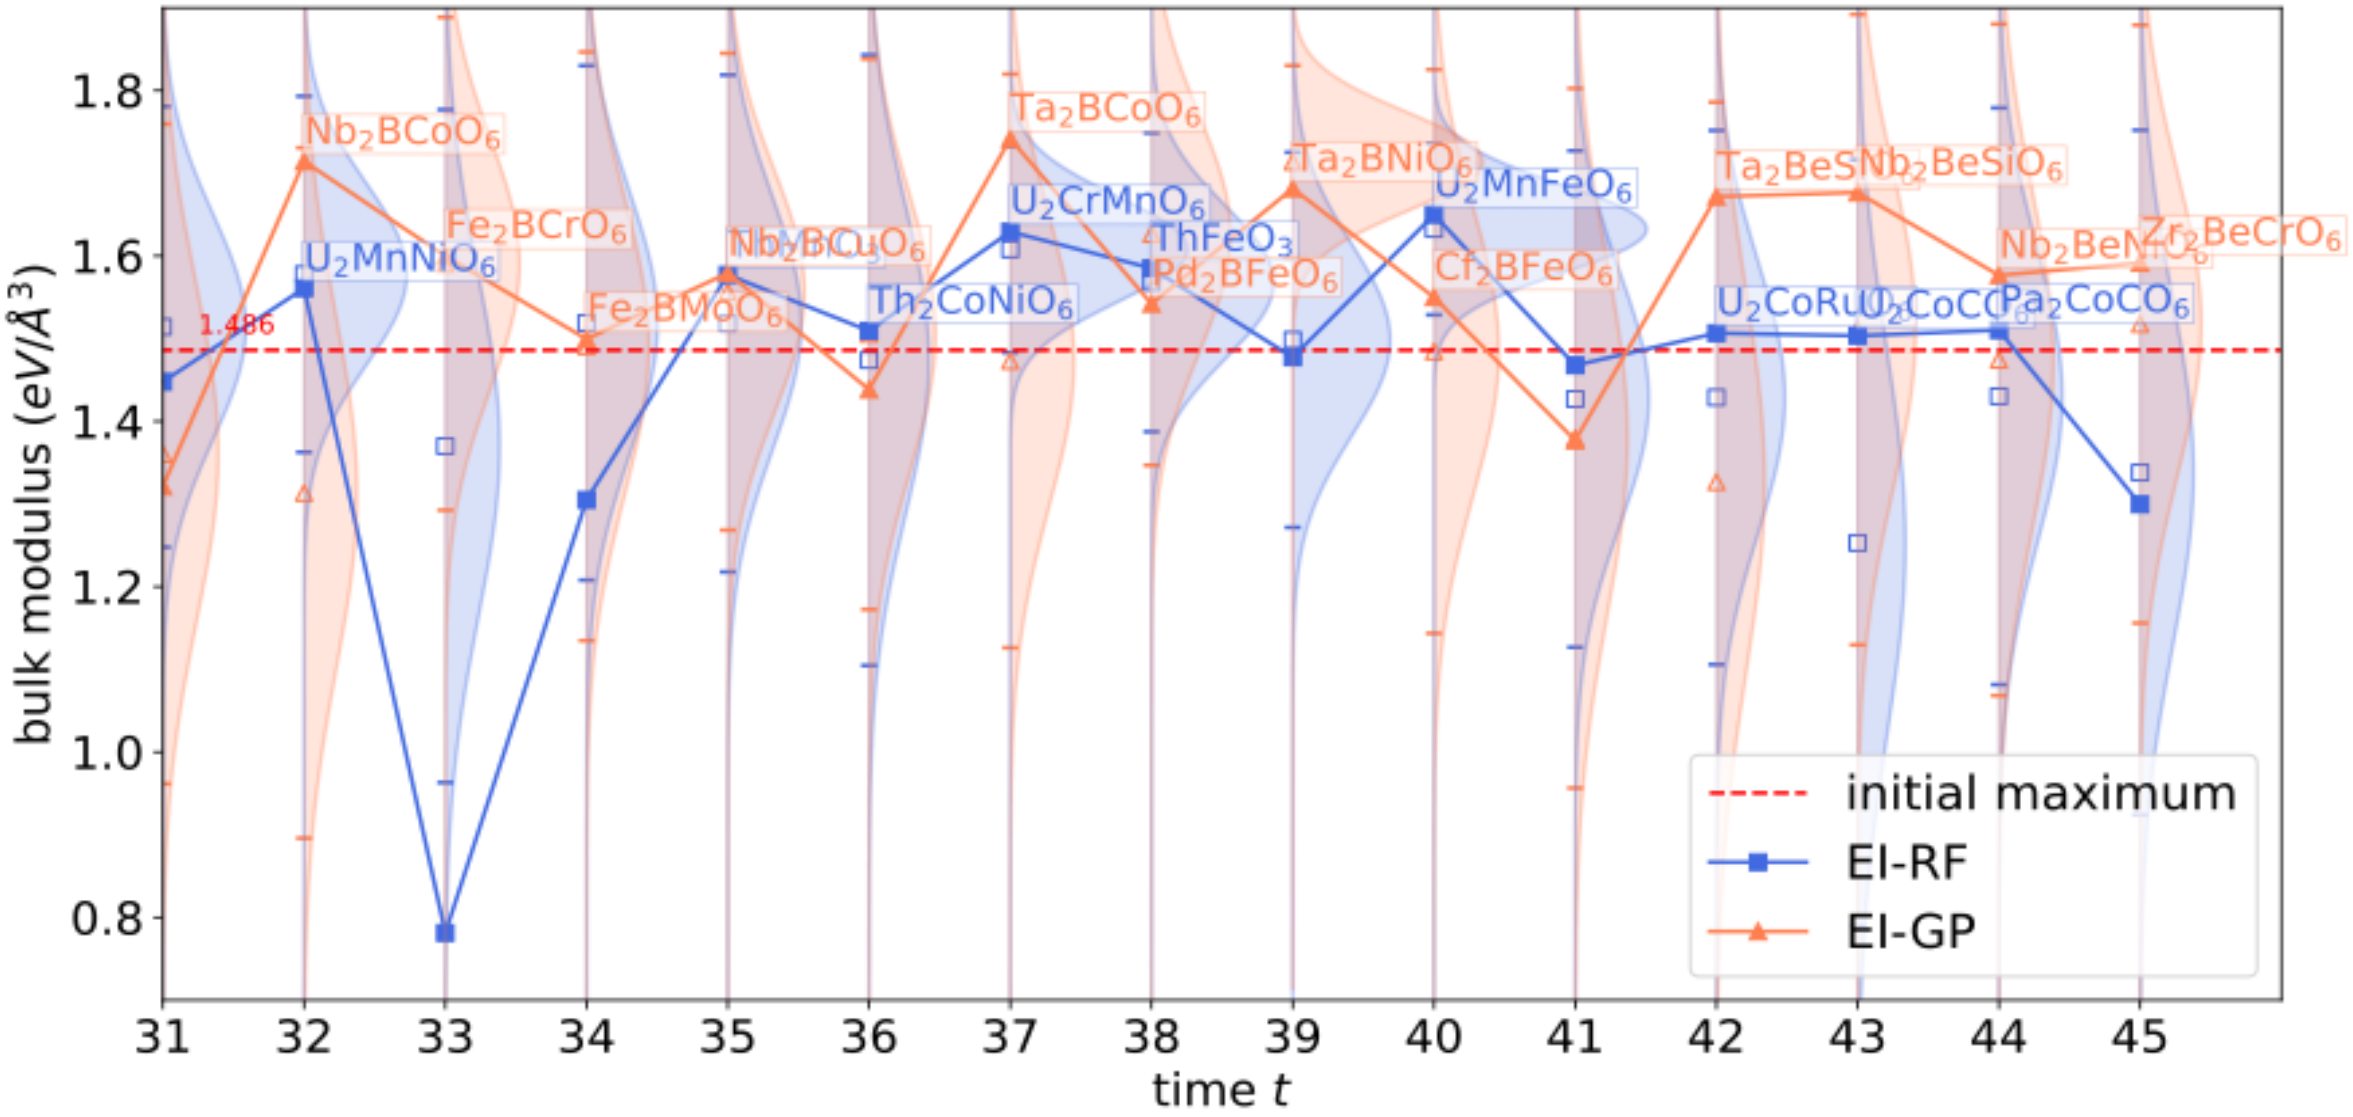

GP outperforms RF due to Uncertainty Quantification

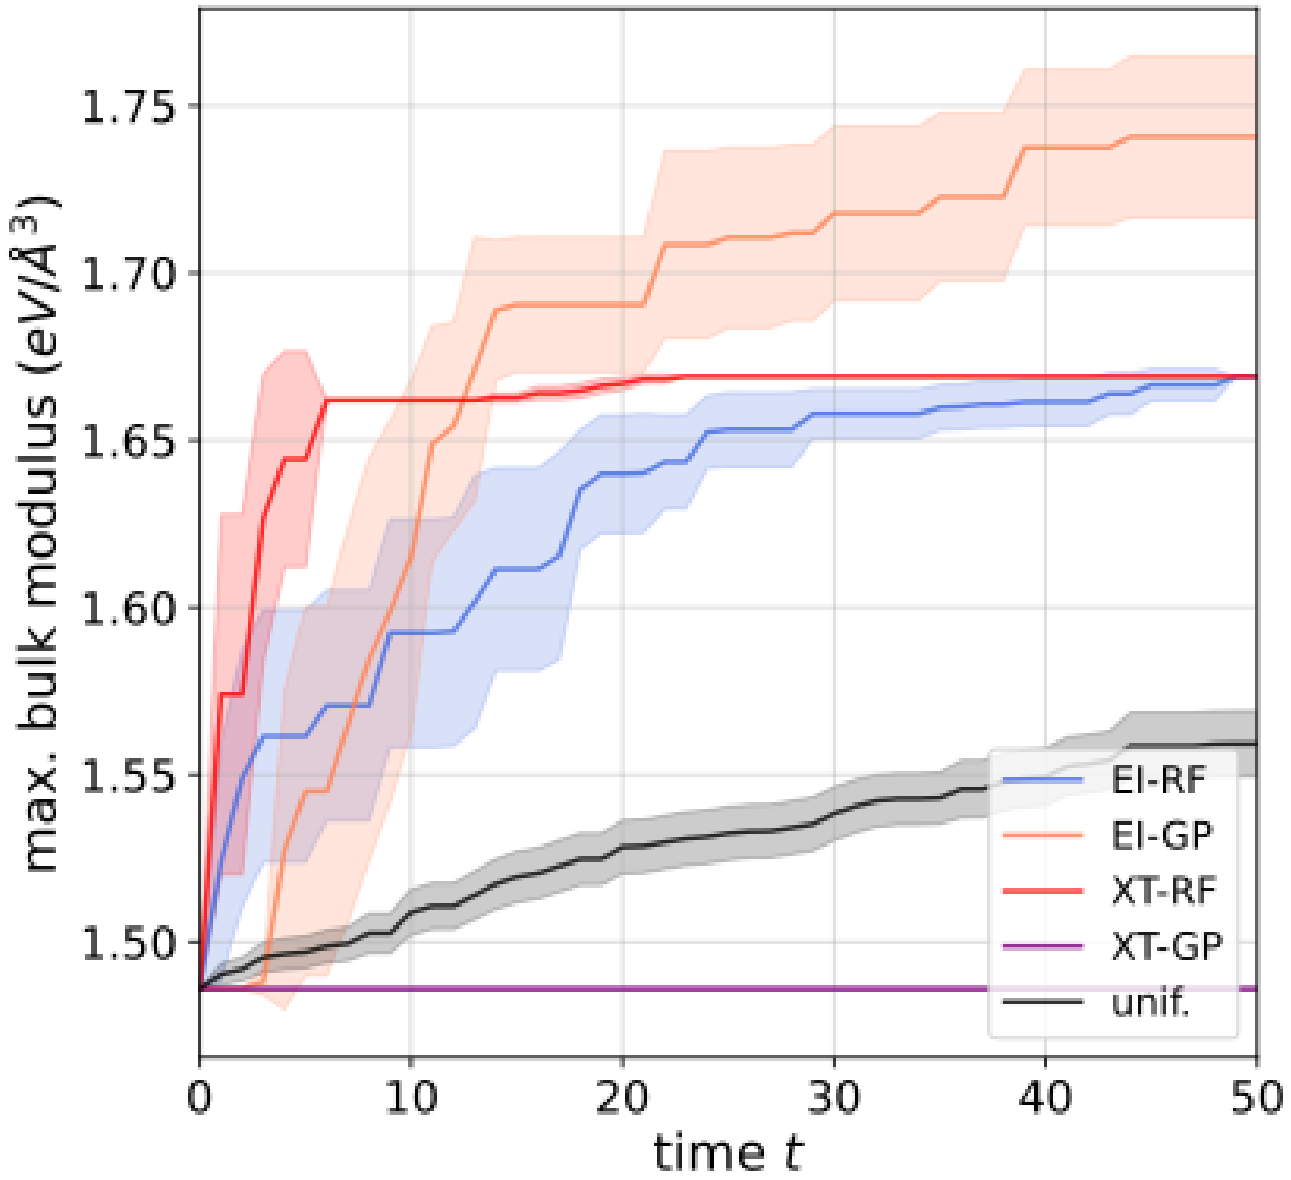

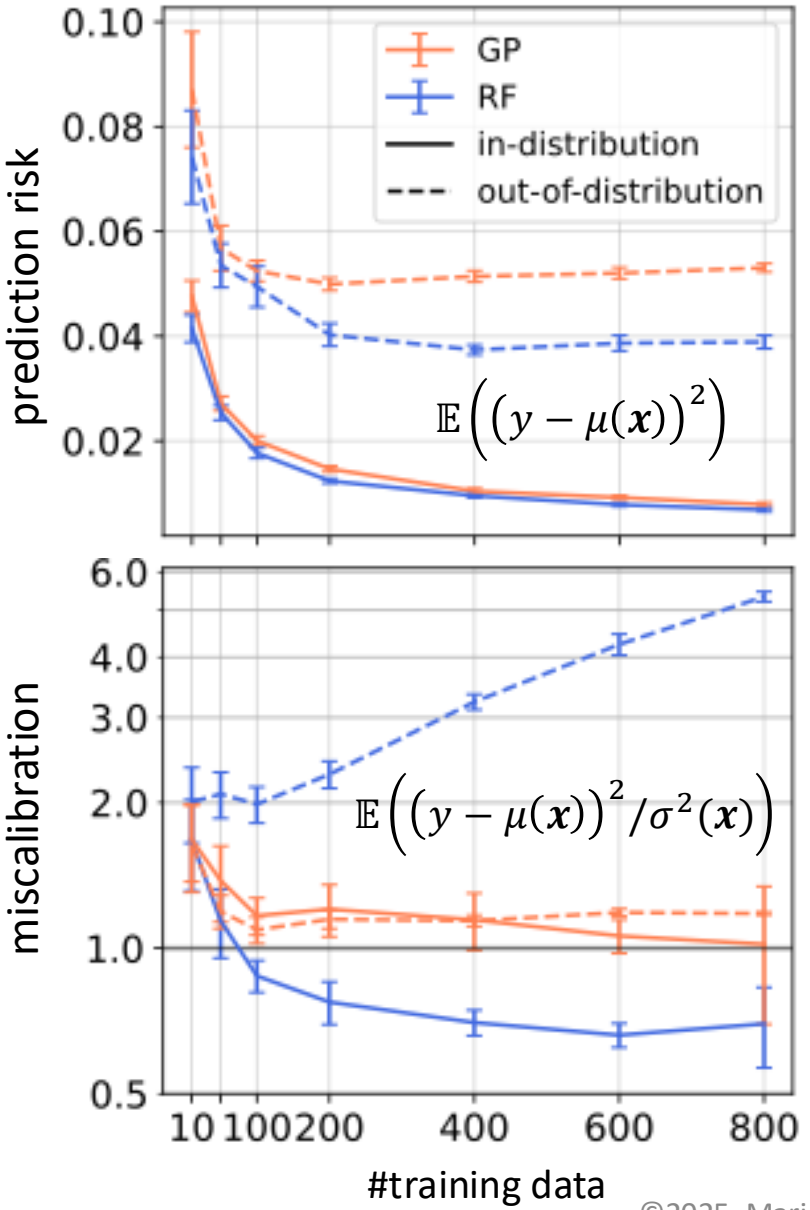

Outline

1. Personas of Agentic AI
2. *Explorer*: mapping polymerisation-induced self-assembly phases
3. *Engineer*: optimising double perovskite hardness
4. *Collector*: discovering novel metal-organic frameworks
5. Problems and directions

Discovering Novel Metal-organic Frameworks

Phuc “Felix” Luong

Zichao Rong

Omar Yaghi

Zihao Chen

Laurent El Ghaoui

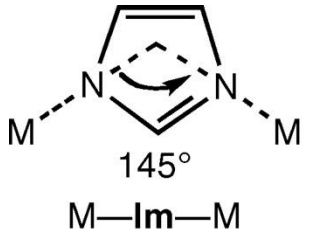

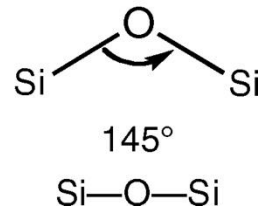

[*Acc. Chem. Res.* 2010, 43, 1, 58-67]

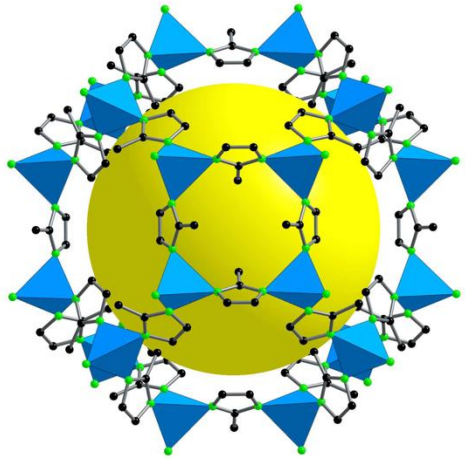

Goal: autonomous materials discovery platform

[Z Rong, Z Chen, P Luong et al. "Algorithmic iterative reticular synthesis of zeolitic imidazolate framework crystals." *Nature Synthesis* (2025): 1-11.]

Collectors Maximise Number of Distinct Discoveries

Ideas

- Split inputs $x = (z, v)$ into "setup" z and "recipe" v
- Choose desired outcomes $Y_+ \subset Y$
- reward is unique combinations of setups and desired outcomes realized by a recipe

$$r_t = \#\{(z, y) \in Z \times Y_+ : (z, v, y) \in s_t \setminus s_{t-1}\}$$

Example: ZIF synthesis

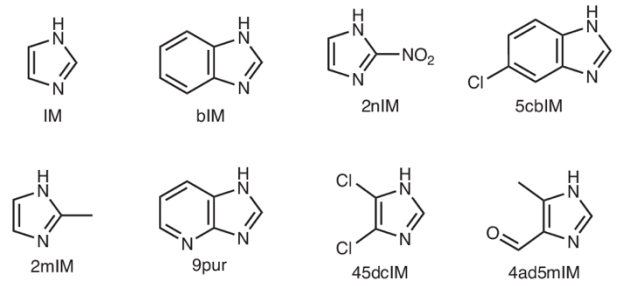

Z precursors,
e.g., $Z = \{\mathrm{Im}, \mathrm{bIM}, 2\mathrm{mIM}, \ldots\}$

V experimental conditions,
e.g., $\mathrm{Conc} \times \mathrm{Temp} \times \mathrm{Time}$

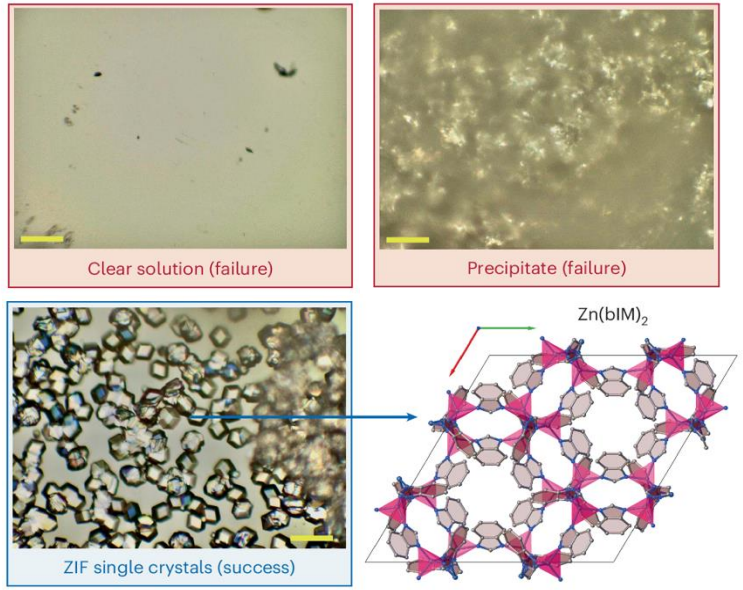

Y_+ targeted crystal structures
e.g., $\mathrm{Y_+} = \mathrm{ZIF\ Crystals}$

Challenge: Experiment Batching

Single experiment expected reward:
unique count of new setup / succ. outcomes

$$\mathbb{E}(r \mid z, v) = \begin{cases} p(y \in Y_+ \mid z, v) & \text{if } \nexists v, (z, v, y) \in s \\ 0 & \text{otherwise} \end{cases}$$

Batch actions $\boldsymbol{a} = (x_1, \ldots x_k) \in X^k$
…describe k experiments $x_i = (z_i, v_i)$

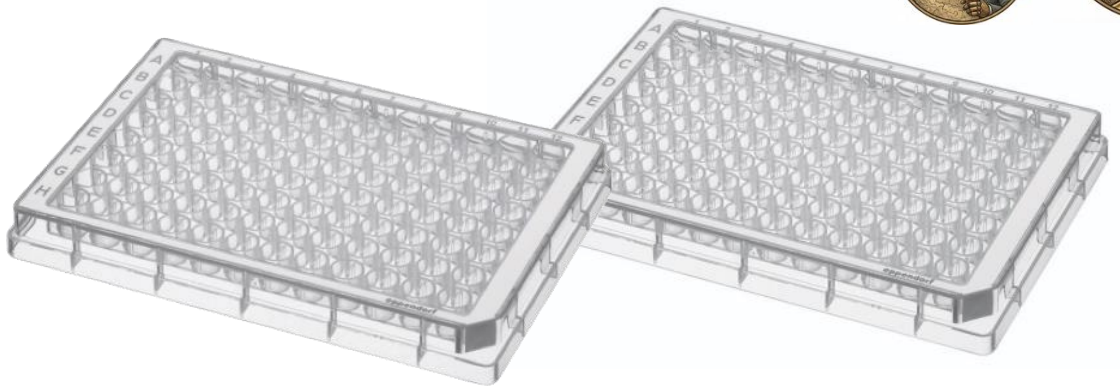

microplates for high-throughput experimentation

Expected batch reward:

$$\mathbb{E}(r \mid \boldsymbol{a}) = \sum_{(z,y) \notin s} \mathbb{P}\left(\bigcup_{i:z_i=z} \{y_i \neq y\} \mid \boldsymbol{a} \right)$$

Constraint

- split recipe variables $v = (c, o)$
- only c can change per well / slot (*composition*)
- other vars o fixed per plate (*temp*, *time*)

Example: two 12-well plates

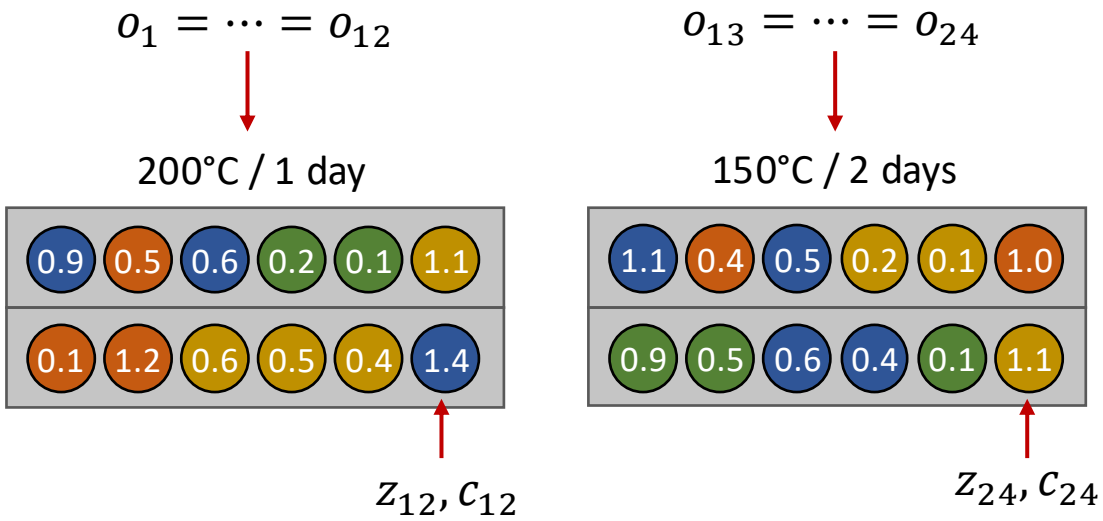

Need Algorithm for Complex Allocation Problem

Batch actions $\boldsymbol{a} = (x_1, \ldots x_k) \in X^k$

…describe s experiments $x_i = (z_i, v_i)$

Expected batch reward

$$\mathbb{E}(r \mid \boldsymbol{a}) = \sum_{(z,y) \notin s} \mathbb{P}\left(\bigcup_{i:z_i=z} \{y_i \neq y\} \mid \boldsymbol{a} \right)$$

precursor / outcome

joint success probability $p_{z,y}$

Constraint

- split recipe variables $v = (c, o)$
- only c can change per well / slot (*composition*)
- other vars o fixed per plate (*temp*, *time*)

Greedy (z, c)-optimisation for fixed o

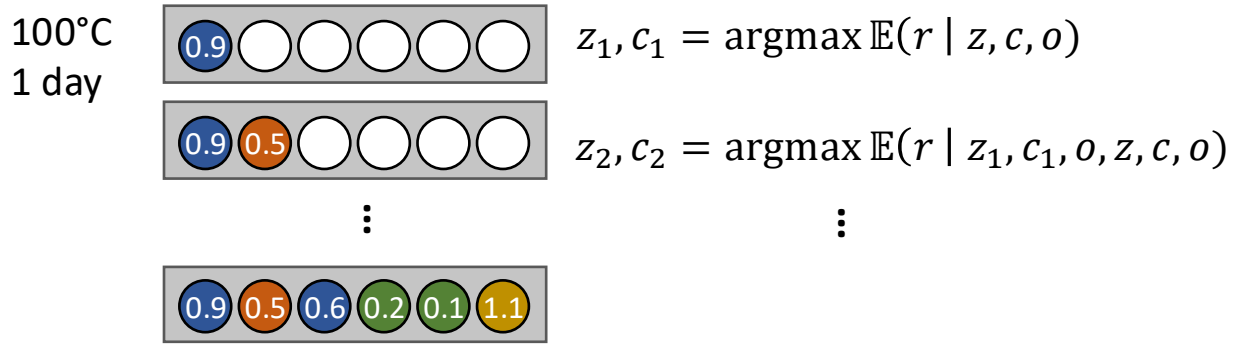

Brute-force condition enumeration

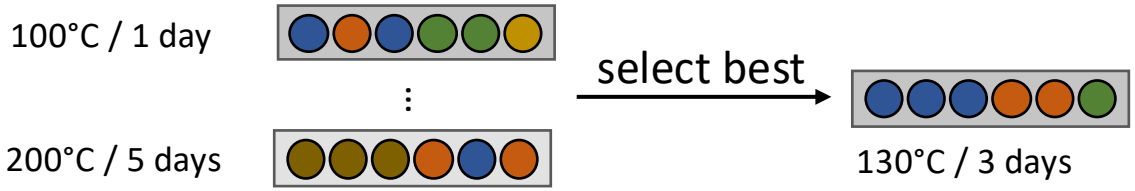

Greedy plate optimisation

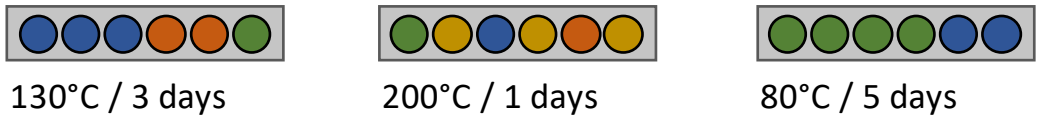

Need to Estimate Joint Success Probabilities

Batch actions $\boldsymbol{a} = (x_1, \dots x_k) \in X^k$

…describe s experiments $x_i = (z_i, u_i)$

Expected batch reward

$$\mathbb{E}(r \mid \boldsymbol{a}) = \sum_{(z,y) \notin s} \mathbb{P}\left(\bigcup_{i:z_i = z} \{y_i \neq y\} \mid \boldsymbol{a} \right)$$

precursor / outcome

joint success probability $p_{z,y}$

Can express success through failure probabilities

$$p_{z,y} = 1 - \prod_{i:z_i = z} q_{z,y}^{(i)}$$

$$q_{z,y}^{(i)} = \mathbb{P}\{y_i \neq y\} / \mathbb{P} \bigcap_{\substack{j:z_j = z \\ i<j}} \{y_j \neq y\}$$

Need to evaluate effect of experiment on batch reward

130°C / 3 days

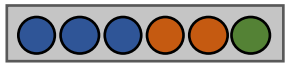

200°C / 1 days

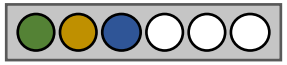

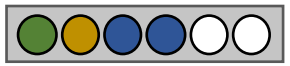

$\mathbb{E}(r \mid x_1, \dots, x_i, z, v)$

$\mathbb{E}(r \mid x_1, \dots, x_i, z', u')$

Idea: inject “hallucinated” failures

a_1	x_1	y_1
	…	…
	x_k	y_k
a_2	x_{k+1}	y_{k+1}
	…	…
	x_{2k}	y_{2k}
	x_{2k+1}	0
	x_{2k+2}	0

state s_2

injected points

$$q_{z,y}^{(i)} = p\left(y \mid z; s_t \cup \left\{(x_j, 0) : z_j = z, j < i\right\}\right)$$

Design Models based on Linker Representation

Random forest classifier

$$y \mid \boldsymbol{x} \sim \mathrm{Bern}\left(\mathbb{E}\left(f(\varphi(\boldsymbol{x}))\right)\right)$$
$$f \mid \boldsymbol{y}_1 \sim \mathrm{GreedyTree}(\tilde{\boldsymbol{y}}_1, \tilde{\boldsymbol{X}}_1)$$

Gaussian process classifier

$$\boldsymbol{f} \sim \mathrm{N}_{|X|}(\boldsymbol{0}, \boldsymbol{K}) \rightarrow k_{i,j} = \exp\left(-\left\|\varphi(x_i) - \varphi(x_j)\right\|^2 / 2\alpha\right)$$
$$y \mid f \sim \mathrm{Bern}\left(\sigma(f)\right)$$
$$\boldsymbol{f}_2 \mid \boldsymbol{y}_1 \sim \mathrm{N}_{|X_2|}\left(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{2,2}\right)$$
$$\boldsymbol{K}_{2,1}\boldsymbol{K}_{1,1}^{-1}\hat{\boldsymbol{f}}_1 \qquad \boldsymbol{K}_{2,2} - \boldsymbol{K}_{2,1}\left(\boldsymbol{K}_{1,1} + \boldsymbol{W}^{-1}\right)^{-1}\boldsymbol{K}_{1,2}$$

Features:

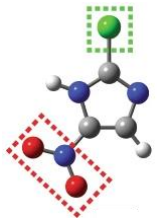

Functional group counts

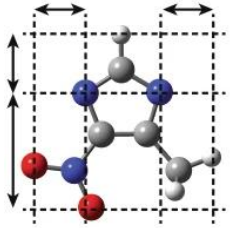

Linker dimensions

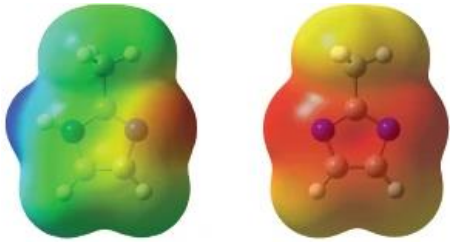

QM properties of neutral/deprotonated linker

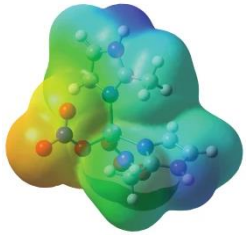

QM properties of pre-nucleation species

RFC-Collector retains Advantage in Batch Mode

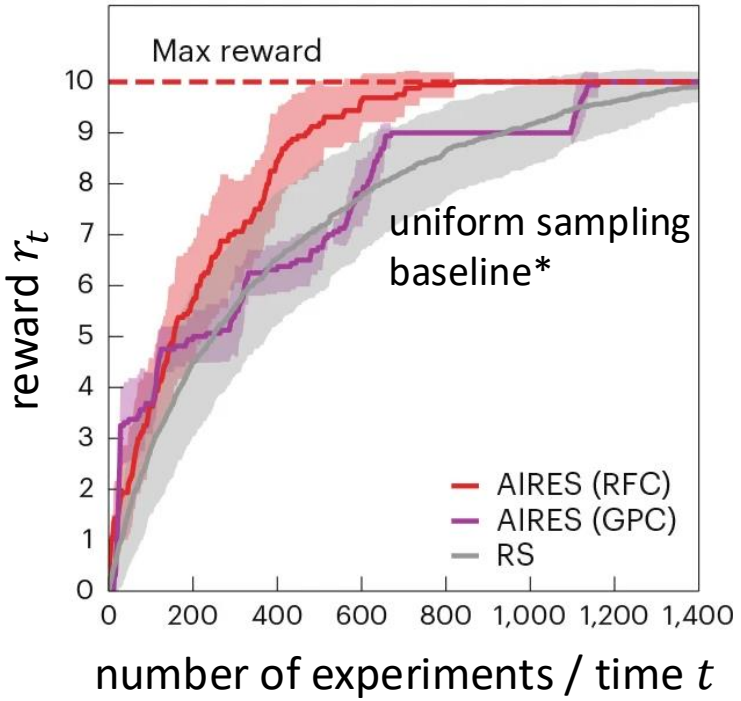

Sequential

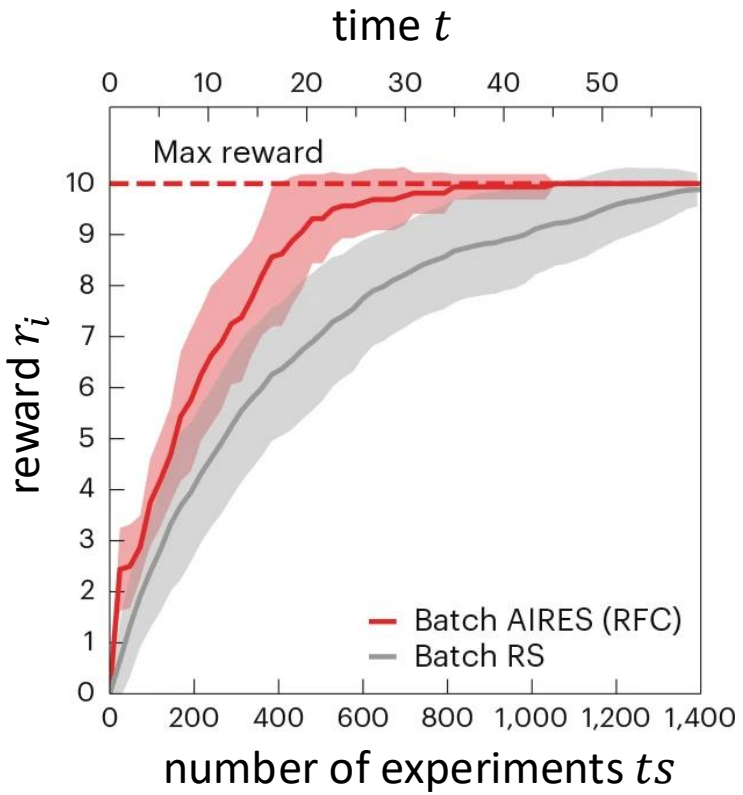

Batch (2 plates, 12 wells each)

$*a_t \sim \mathrm{Unif}\{(z,v) : \exists y \in \mathrm{Y}_+, \nexists v', (z,v',y) \in s_t\}$

Strategy Drastically Shrinks Tails of Discovery Times

Time to discovery of recipe leading to outcome $y \in Y_+$ for setup z:

$$t(z, y) = \min\{t \in \mathbb{N} : \exists v, (z, v, y) \in s_t\}$$

Depends on required recipe specificity:

$$\mathbb{E}_{v \sim \mathrm{Unif}(V)}\, p(y_+ \mid z, v)$$

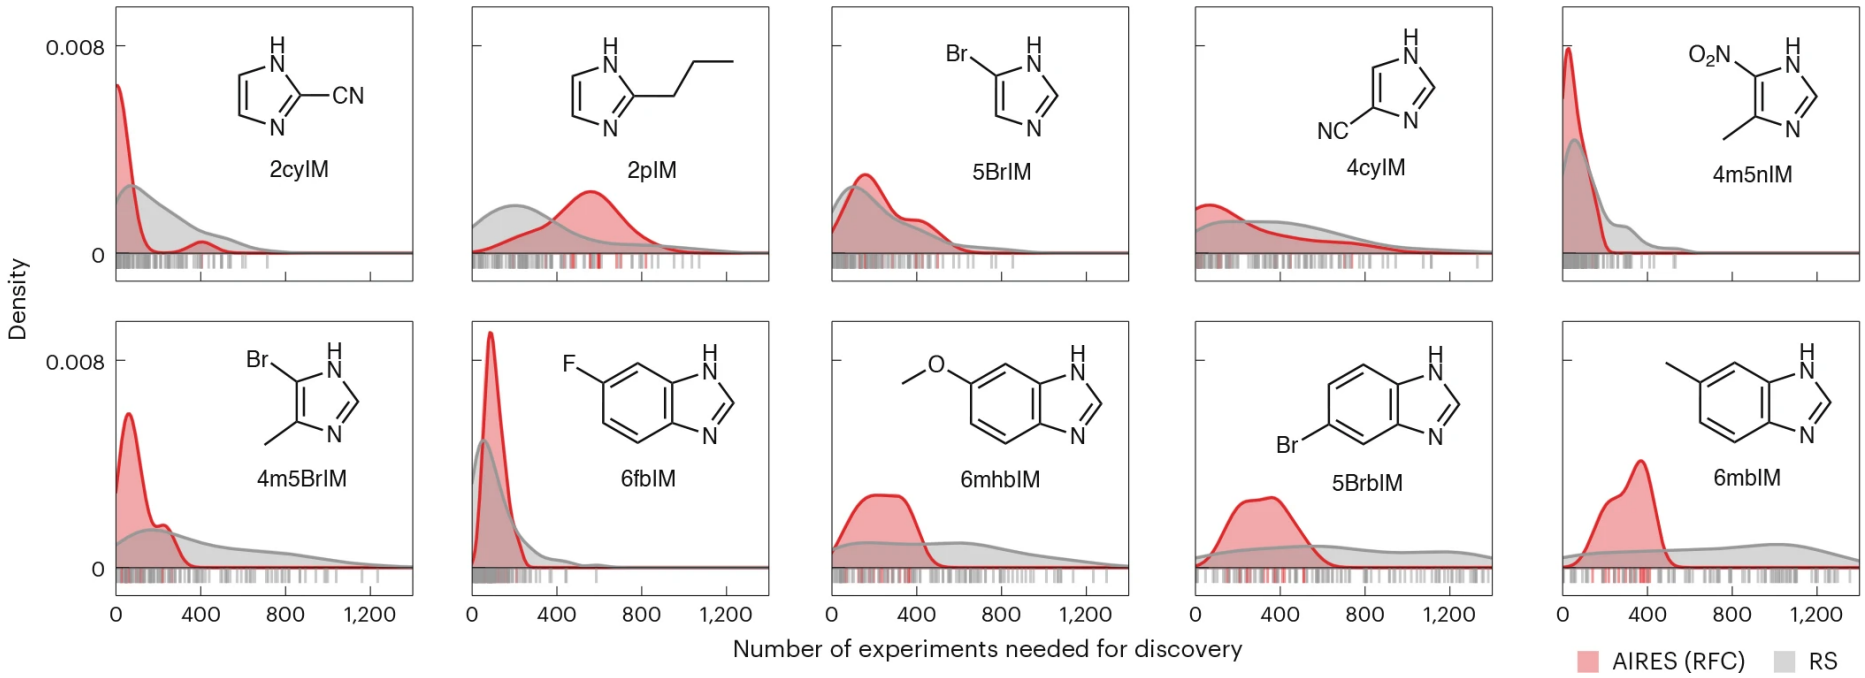

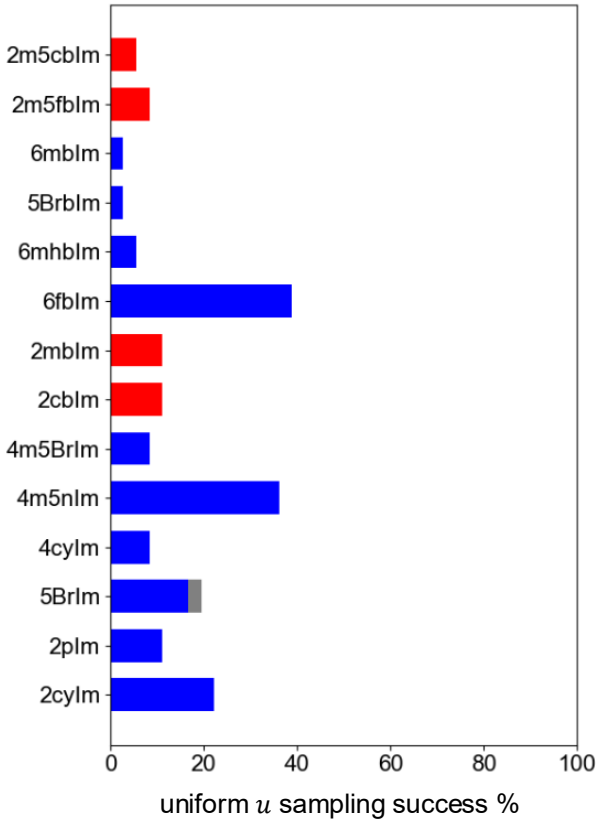

Outline

1. Personas of Agentic AI
2. *Explorer*: mapping polymerisation-induced self-assembly phases
3. *Engineer*: optimising double perovskite hardness
4. *Collector*: discovering novel metal-organic frameworks
5. Problems and directions

Directions to Explore in Collaborations

1. Interpretability: rule-based models
2. Representation Learning: GNNs, LLMs, pre-training
3. Uncertainty quantification
4. Grid-free optimisation
5. Experiment look-ahead and batching