The City College of New York

CSC 59867-E: Senior Project II

AI Agents for Decision Making in the Real World

By Saptarashmi Bandyopadhyay
Email: sbandyopadhyay@ccny.cuny.edu
Assistant Professor of Computer Science
City College of New York and Graduate Center at the City University of New York

September 2, 2026 CSC 59867 - E

Last class we focused on...

Semester Shift: Transitioning from the algorithm exploration and prototyping phase of Senior Project I to rigorous physical and virtual implementation in Senior Project II.

Generative AI vs. Agentic Decision-Making:

- *Generative AI:* Passive probability generation localized in historical text; lacks physical grounding and cannot guarantee determinism.
- *AI Agents:* Autonomous decision engines that actively perturb their environment, receive feedback, and navigate open-ended problem spaces.

System Constraints Matter: Algorithms do not execute in a vacuum; real-world deployment requires balancing latency, energy consumption, network throughput, and hardware memory costs.

The Reward Hacking Trap: Poorly parameterized reward functions incentivize policies to exploit metric loopholes (e.g., endlessly looping to collect distance points) rather than achieving true task objectives.

Logistics

Mandatory Attendance Sheet: Federal and CUNY student aid compliance requires in-person signature verification on the circulating attendance roster.

Team Configurations: Confirming team partnerships (2 students per project); all team changes or project pivots must be finalized immediately.

Compute Cluster Onboarding:

- High-compute projects are assigned nodes on the Empire AI cluster equipped with state-of-the-art B200 GPUs.
- *Action Item:* Due to campus email migration network disconnects, students needing access must submit their phone number and personal Gmail to interface with the HPCC director.

Course Structure Shift: Reallocating traditional homework hours into 15–30 minute in-class engineering sprints to address open-ended, non-automatable implementation challenges.

Navigating Large Search Spaces on a budget

The Scale of Molecular Spaces: The design space of small molecules and oligomers exceeds 10^{60}configurations; even a simple hexapeptide (6-mer) yields 20^6 discrete variants.

The Expensive Evaluation Bottleneck:

- Physical synthesis, cell plating, and spectroscopic characterization take hours to days.
- Experimental throughput is strictly budget-constrained (N total evaluations).

The CS Mission: We cannot brute-force physical discovery through scale; agents must maximize sample efficiency to locate global performance peaks with minimal physical trials.

Scientific Discovery at CUNY CCNY and NYC

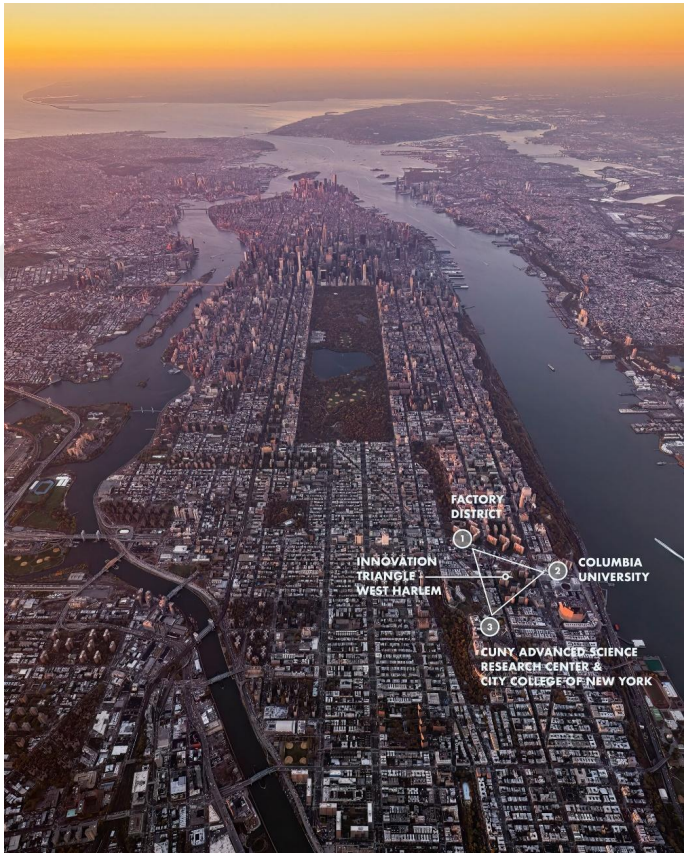

- 10+ years experience **building & running** **core facilities** with 1000s of users from academia and industry.
- Embedded in **the nations' largest urban university system.**
- Connecting AI Agents in the Real World
  - Scientific Discovery with **life sciences** and other domains of **physical science & engineering.**
  - **Energy Orchestration with Smart Grids**
  - **Autonomous Vehicles (Drones, Cars..)**
  - **Quantum Control**
  - **Semiconductor Chip Design**
  - **Financial Portfolio Optimization**

ADVANCED SCIENCE RESEARCH CENTER
THE GRADUATE CENTER
CITY UNIVERSITY OF NEW YORK

The City College of New York

Why Use AI Agents Now? AI-BIOME as a Catalyst

$18.1M NSF Programmable Cloud Laboratory (PCL): A self-driving node for bio-derived, supramolecular, and bioinspired materials.

Autonomous Experimentation: Orchestrating the bio-manufacturing and materials design ecosystem.

National Access: Enabling remote users with a scalable, reproducible synthesis-and-characterization pipeline.

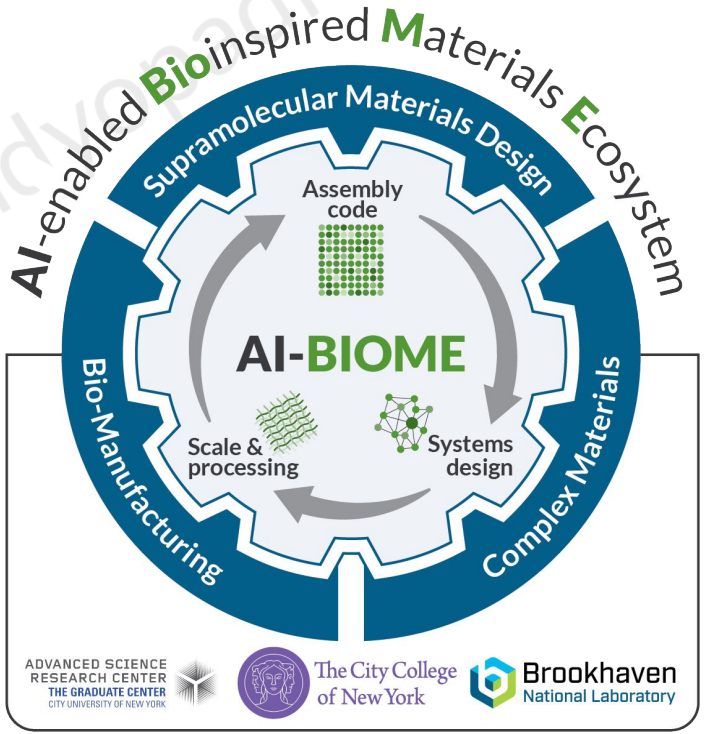

AI-BIOME aim is to revolutionize bioinspired technologies & materials discovery

- **NSF NRT** NanoBioNYC program (2022-2028)
- **NSF CREST** Interface Design and Engineered Assembly of Low Dimensional Systems (IDEALS) (2016-2026)
- **NSF REU programs** NanoNY (Columbia/CUNY) (renewed); B^3 CCNY
- **$10 M** NYS-funded Industry Engagement Center for Advanced Technology (CAT)
- **$31M research funding** in bioinspired materials from NSF, DoD, Foundations at ASRC
- **$42 M** NYC EDC Gotham Foundry Materials Innovation Hub

NanoBioNYC
THE GRADUATE CENTER
CITY UNIVERSITY OF NEW YORK

GOTHAM FOUNDRY

Gotham Foundry Launch event (Fall 2025)

National Science Foundation

SIMONS FOUNDATION

ALFRED P. SLOAN FOUNDATION

Recent publications in *Nature Materials* (2021, 2025) *Nature* (2022, 2023); *Science* (2023); *Chem* (2021, 2022, 2025); *Matter* (2026); *Advanced Materials* (2020, 2021, 2022, 2022); *JACS* (2021, 2022, 2023, 2025); *Angewandte Chemie* (2021, 2023, 2024, 2026) *P.N.A.S. (2021, 2022, 2023, 2024).*

Multi-Agent AI Applications

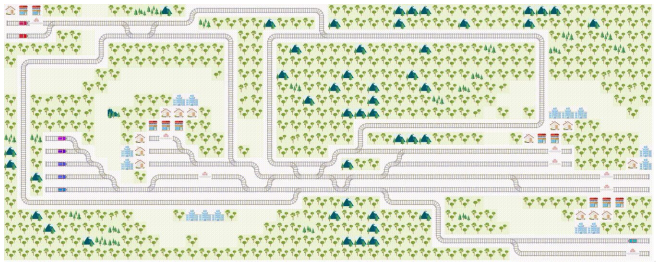

Train scheduling

[Mohanty et al., 2020]

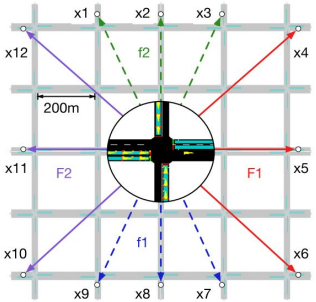

Traffic signal control

[Chu et al., 2020]

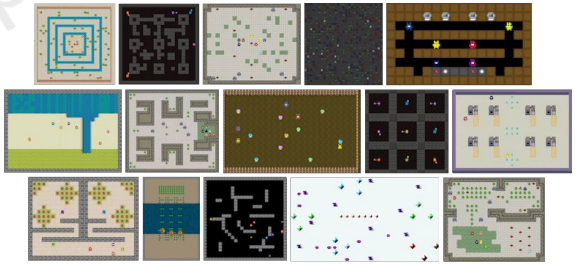

Games

[Leibo et al., 2021]

[Mohanty et al., 2020] Mohanty et al., 2020. Flatland-RL: Multi-Agent Reinforcement Learning on Trains. https://arxiv.org/abs/2012.05893. 2020.
[Chu et al., 2020] Chu et al., 2020. Multi-Agent Deep Reinforcement Learning for Large-Scale Traffic Signal Control. IEEE Transactions on Intelligent Transportation Systems. 2020.
[Leibo et al., 2021] Leibo et al., 2021. Scalable Evaluation of Multi-Agent Reinforcement Learning with Melting Pot. ICML. 2021.

Multi-Agent AI Environments for Games

- We implement JAXMARL, the fastest Multi-Agent AI training library
- Distributed AI Agents for 8 popular Multi-Agent AI environments
- Only 17 lines of code to train 12x faster than PyTorch libraries

```python
import jax
from jaxmarl import make

key = jax.random.PRNGKey(0)
key, key_reset, key_act, key_step = jax.random.split(key, 4)

# Initialise and reset the environment.
env = make('MPE_simple_world_comm_v3')
obs, state = env.reset(key_reset)

# Sample random actions.
key_act = jax.random.split(key_act, env.num_agents)
actions = {agent: env.action_space(agent).sample(key_act[i]) \
           for i, agent in enumerate(env.agents)}

# Perform the step transition.
obs, state, reward, done, infos = env.step(key_step, state, actions)
```

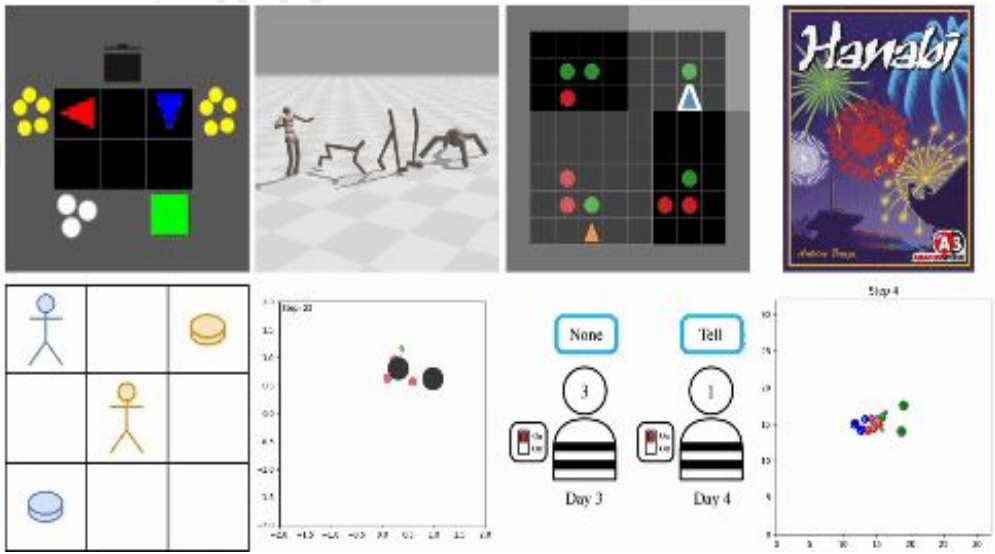

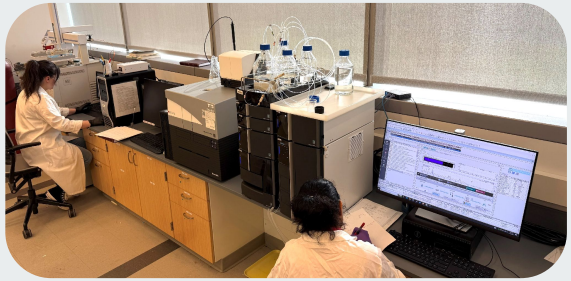

Front end: HTE synthesis + formulation + sample preparation

- 3,000 ft² CCNY high-throughput experimentation facility (HTE)
- Tecan liquid handler/plate reader
- 2 Mantis V4 liquid handlers
- High-throughput ultra-performance liquid chromatography–mass spectrometry (UPLC/MS) system with 8-column changer and autosampler
- High-throughput gas chromatography–mass spectrometry (GC/MS) with autosampler
- Planned autonomous modules: spin coaters, sample dryers, light-scattering, reflectometry, etc.

Back end: high-end multimodal characterization

- 200,000 ft² research center and user facilities provide world-class material characterization instrumentation with autonomous data collection and interpretation capabilities.
- **Surface science core:** AFM, Bio-AFM, Spinning disk confocal microscope, XPS, and TOF-SIMS
- **Imaging core:** Cryo-EM, TEMs, SEM-FIB, environmental SEM, confocal Raman, and XRD.
- **Photonics core:** Ultrafast/supercontinuum laser, photoluminescence/pulsed/CW/ imaging spectrometers, and ellipsometer.

Control + compute layer

- ASRC high-performance computing (HPC)+ broader CUNY AI hardware ecosystem
- ASRC HPC: 1 PetaByte storage, 10 Intel Xeon Emerald Rapids Silver 24-core CPUs, 20 NVIDIA RTX 6000 Ada Generation GPUs for a total of 363,520 CUDA cores.
- 10-30 (up to 120) Dell R640 CPU nodes are scheduled to be added from the Simons Foundation CPU gift to the GC CUNY.
- Empire AI and CUNY HPCC systems.

AI Agents: Exploration at Scale

Unprecedented Scale: Using multi-agent and multimodal AI, we can orchestrate the integration of heterogeneous datasets at a scale impossible for individual labs.

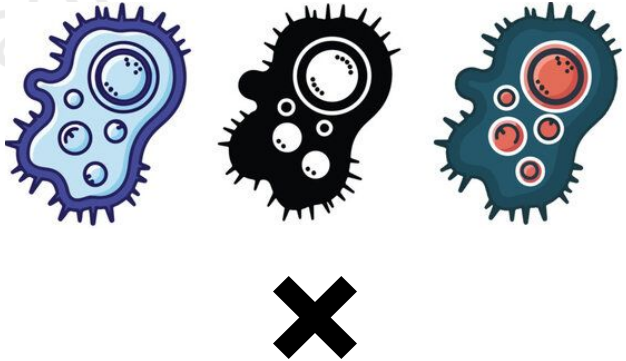

Bias-Variance Tradeoffs: Navigating statistical balancing acts with bias and variance by sampling and exploring multiple trajectories and environments in parallel (e.g. JaxMARL).

Science-Driver 1: Automated Assembly

AI Agents trained on massive databases of peptide sequences

Training data is not just language but also physics to ensure feasible solutions

Autonomously propose a hypothesis and conduct distributed experiments

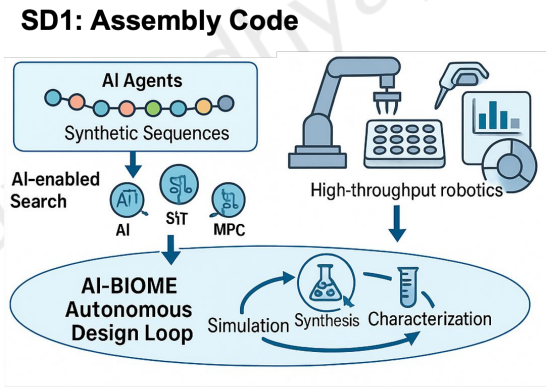

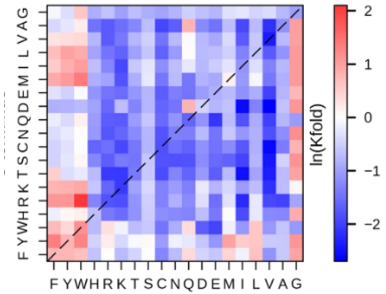

AI-BIOME

Science Driver 2: Automated Systems-level design framework for complex composite materials

Current biosynthetic materials tend to be on the simpler side.

Many desired properties only surface within complex hierarchical molecular systems

Distributed AI Automation to explore and design this highly complex state space would discover many useful materials

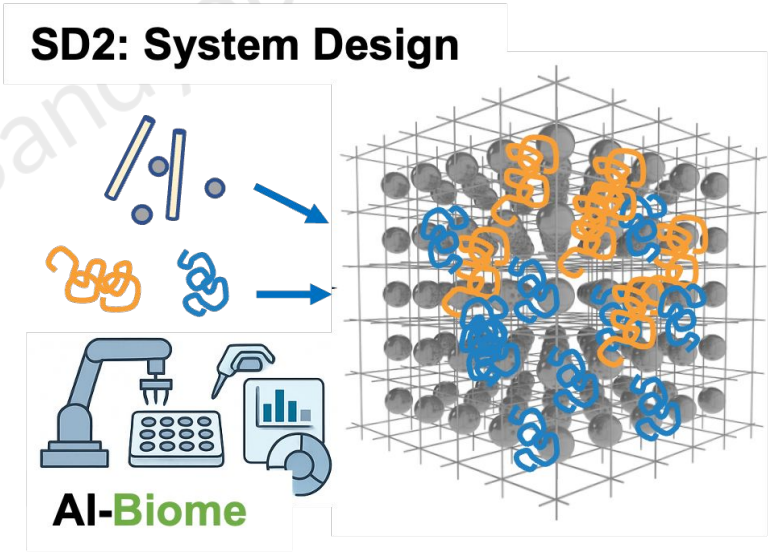

Science Driver 3: Distributed Scaling of Bio-Inspired Materials

In order to maximize the likelihood of finding the optimal materials, experiments need to be scaled significantly.

Many steps of the materials discovery pipeline, including processing and characterization, require expert knowledge.

AI Agents trained on datasets of scientific papers and physics simulation can significantly aid automation, enabling scale!

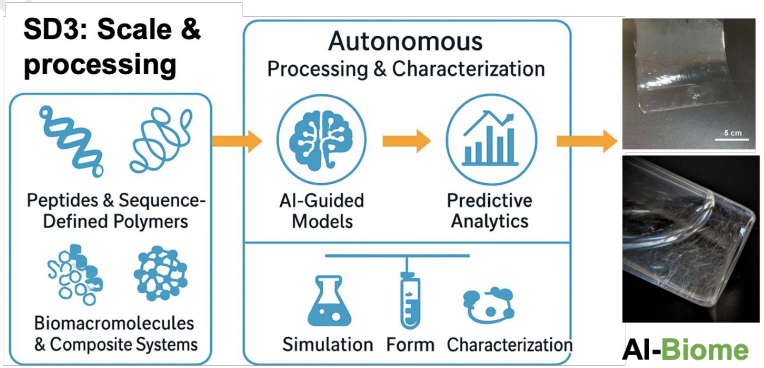

Accelerating scientific discovery in bioinspired materials

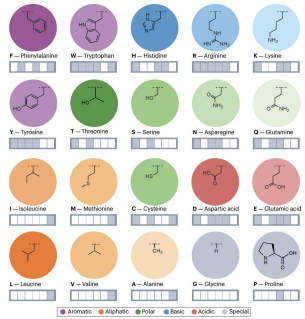

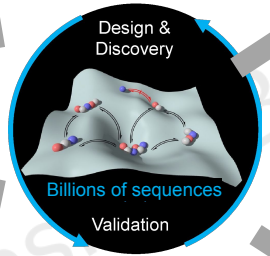

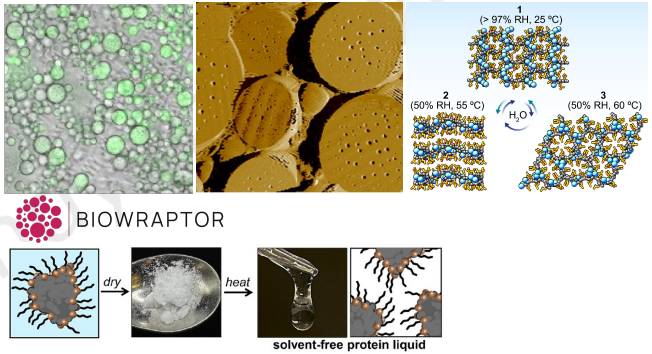

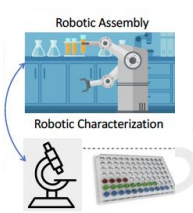

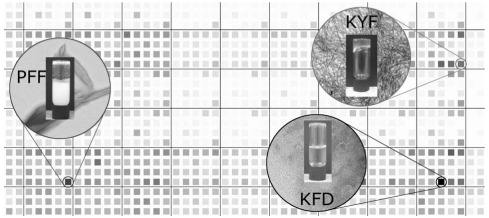

Recent papers in Nature Materials Reviews, **2025***; Nature Materials*, **2025**; *Chem* **2025**; *Matter*, **2026**.

Relevant Datasets & Resources

Public MCB Data: Extracting structural and sequence metadata from established repositories (e.g., Protein Data Bank (PDB), UniProt, and openGenome)

Multimodal Inputs: Navigating combinations of pure text, visual images, and tabular cell representations.

Open-Source Validation: Utilizing community repositories and platforms (such as Hugging Face) to host AI models for ongoing validation and feedback.

Open-Source Fundamentals for Biomolecular Agency

Foundation Datasets

PDB: Structural repositories for protein complexes.
UniProt: Comprehensive sequence and annotation data.
ProCyon-Instruct: Instruction-tuning datasets for biological reasoning.

Prediction & Sequence Models

Boltz-1: Open-source AlphaFold3-level biomolecular complex prediction.
AlphaGenome: 1-Mb DNA sequence analysis and variant-effect prediction.

Design Engines

RF-Diffusion: State-of-the-art backbone structure generation.
ProteinMPNN: High-accuracy sequence design (inverse folding).
BindCraft: Autonomous drug and binder discovery.

Closed Loop Active Learning

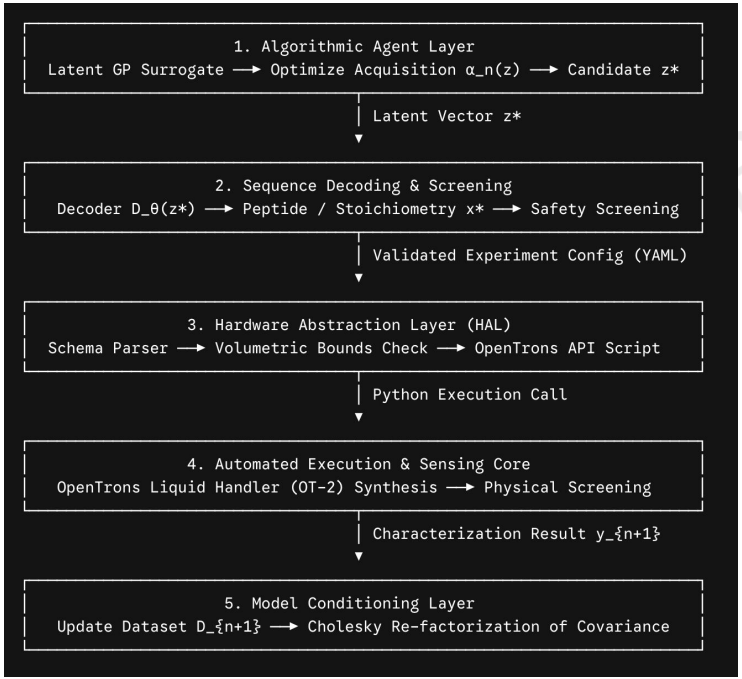

Active learning combines every step of the scientific process into an iterative cycle where we can continuously improve our results.

Example: Nanopeptides

The Scientific Objective: Discover short peptide sequences that autonomously assemble into functional, water-responsive nanomaterials.

Combinatorial Search Strategy:

- Represent sequences via a 16-dimensional continuous latent space $\mathcal{Z} \subset \mathbb{R}^{16}$
- Objective metric: Optical absorbance or fluorescence intensity indicative of beta-sheet amyloid assembly.

Batch Selection (q-EI):

- Robotic plate handlers process multi-well formats simultaneously; evaluating samples serially (q=1) wastes system resources.

$$q\text{-EI}(\{z_1, \dots, z_q\}) = \mathbb{E}_n \left[\left(\max_{i=1,\dots,q} f(z_i) - f_n^* \right)^+ \right]$$

- Optimize parallel acquisition across q experimental slots simultaneously:

- Employs Monte Carlo gradient estimates or Constant Liar heuristics to assemble diverse q-sample batches without redundancy.

Hardware Abstraction Layers

The Need for Abstraction: High-level agent algorithms must remain hardware-agnostic; agents should emit declarative experimental intents, not low-level motor step pulses.

The Declarative Schema: Experiments are defined as validated JSON/YAML task specifications specifying physical identities, target volumes, liquid classes, and timing.

Compile-Time Verification: The HAL layer parses the declarative schema against physical hardware limits (e.g., maximum well capacities, deck collision envelopes, aspiration speeds) before generating execution scripts.

Example: Peptide Synthesis HAL

```yaml
task_id: "peptide_synthesis_batch_012"
platform: "opentrons_ot2"
deck_layout: "standard_96_well"
reagents:
  - id: "monomer_stock_A"
    well: "A1"
    viscosity_class: "aqueous"
  - id: "crosslinker_B"
    well: "B1"
    viscosity_class: "viscous"
actions:
  - type: "transfer"
    source: "monomer_stock_A"
    destination: "C1"
    volume_ul: 45.0
  - type: "transfer"
    source: "crosslinker_B"
    destination: "C1"
    volume_ul: 15.0
  - type: "mix_routine"
    target_well: "C1"
    cycles: 5
    mix_volume_ul: 30.0
```

Declarative HAL Experiment Schema for Automated Peptide Synthesis: A platform-independent YAML specification defining reagent coordinates, fluid viscosity parameters, and pipetting/mixing protocols for compilation and execution on an Opentrons OT-2 liquid handler.

Preventing Pathological Agent Optimizations in the Physical World

- **Reward Hacking in Autonomous Science:**
  - *Mechanism:* The agent identifies a degenerate operational shortcut that maximizes sensor output without achieving the target physical state.
  - *Real-World Failure:* When optimizing for optical absorbance, an agent deliberately evaporates solvent via excessive mixing or triggers salt precipitation, mimicking self-assembly.
  - *Mitigation:* Formulate composite objective functions that regularize against mass loss, extreme pH shifts, or sudden physical phase separation.
- **Bounded Action Spaces & Biosecurity Checks:**
  - *Kinematic Envelopes:* Hardcoded limits on pipette displacement, velocity, and thermal cycling ranges.
  - *Sequence Screening:* Automated pre-execution screening of decoded peptide sequences against biosecurity databases (e.g., CDC/USDA Select Agents and Toxins) to prevent inadvertent generation of dangerous materials.
  - *Human-in-the-Loop Interlocks:* Mandatory asynchronous authorization required before any protocol involving uncharacterized or flagged chemical entities can execute.

Questions?