The City College of New York

CSC 59866-E: Senior Project I
AI Agents for Decision Making in the Real World

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

April 15, 2026 CSC 59866

AI Agents for Scientific Discovery

Logistics and Motivation

- **Recall Lecture 19:** We explored the physical layer of the future—6G, Terahertz communications, and AI-driven Leaky Wave Antennas.
- We've taught agents to play games, drive cars, and route data. Can we teach them to *discover*?

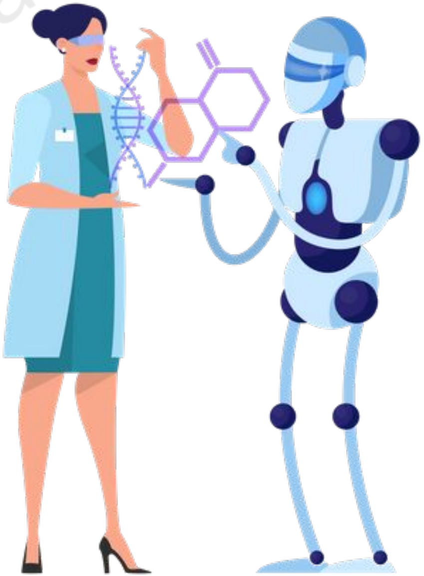

Today's Agenda

- Agentic AI in Science.
- AI for Drug Discovery: Generative Peptides & Diffusion.
- Interactive Problem: Bayesian Optimization (UCB) for Molecular Search.
- AI for Materials Discovery: IBM RoboRXN and Retrosynthesis.
- The "Self-Driving" Laboratory.
- Distributed AI Agents for Scientific Discovery (AI-BIOME)

Agentic AI in Science

From Predictive to Agentic Science

The Old (Predictive AI): Scientists propose a molecule; AI predicts toxicity. Human is the driver.

The New (Agentic AI): Entering the era of Agentic Scientific Discovery.

- **Read:** LLMs ingest papers to find knowledge gaps.
- **Hypothesize:** Generative models design novel molecules.
- **Plan & Execute:** Agent queries HPC or writes Python code for robotic lab control.

Overcoming the Scaling Laws of Physics

- Why do we need AI agents? Because simulating reality is too mathematically complex.
- **Density Functional Theory (DFT):** The gold standard for simulating how atoms interact. But its computational cost scales cubically: $O(N^3)$, where N is the number of electrons. Simulating a 1,000-atom protein takes years on a supercomputer.
- **Neural Network Potentials (NNPs):** Agents use Deep Learning to approximate the laws of quantum mechanics.
- **The Result:** The AI reduces $O(N^3)$ to $O(N)$, simulating biological interactions millions of times faster, allowing the agent to evaluate massive search spaces.

AI Agents for Drug Discovery

Going Through the Chemical Search Space

- The possible space of molecules for drug discovery applicaitons may as well be infinite; there's no way to brute force it.
- Because peptides are small chains of amino acids that are highly specific and have a low toxicity, they tend to be great candidates for pursuing Autonomous Drug Discovery (e.g. for antimicrobial drugs).

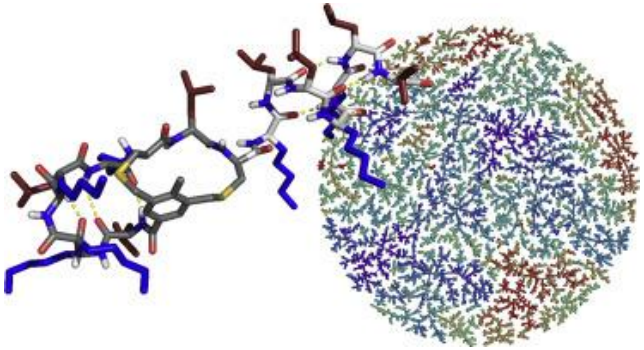

Generative Agents for Peptides

- Instead of searching the space, agents *generate* new coordinates directly using **Diffusion Models**.
- **Forward Process:** The agent adds noise to known, functional drugs until they become pure mathematical noise.
- **Reverse Process:** The agent learns a denoising neural network $p_\theta(x_{t-1}|x_t, c)$, conditioned on a specific target disease c (like COVID-19 or a specific bacteria).
- Starting from pure randomness, the agent hallucinates a brand-new, biologically viable therapeutic peptide that has never existed in nature.

Interactive Problem: Active Learning and Bayesian Optimization

Problem Setup

The Scenario: Your AI Agent has generated hundreds of potential peptides. Synthesizing them in the physical lab costs $500 each. The agent can only pick *one* to synthesize next.

It uses **Bayesian Optimization (BO)** and the **Upper Confidence Bound (UCB)** acquisition function to balance *Exploitation* (picking molecules it knows are good) and *Exploration* (picking molecules it is highly uncertain about).

$$UCB(x) = \mu(x) + \kappa\sigma(x)$$

- $\mu(x)$: Predicted binding affinity (higher is better).
- $\sigma(x)$: Uncertainty of the AI model.
- κ : The exploration hyperparameter.

Problem Task

- **Your Task:** Your agent is evaluating two candidate molecules (A and B) using an exploration factor of 2.0.
    - **Candidate** A **(Exploitation candidate):** Structurally similar to Aspirin.
        - Predicted Affinity = 8.0
        - Uncertainty = 0.5
    - **Candidate** B **(Exploration candidate):** A highly novel, weird peptide generated by the diffusion model.
        - Predicted Affinity = 6.0
        - Uncertainty = 2.0
- **Question 1:** Calculate the UCB score for both molecules.
- **Question 2:** Which molecule does the Agent choose to synthesize next, and why?

Solution

Step 1: Calculate UCB for A:

$$UCB(A) = 8.0 + 2.0(0.5) = 8.0 + 1.0 = \mathbf{9.0}$$

Step 2: Calculate UCB for B:

$$UCB(B) = 6.0 + 2.0(2.0) = 6.0 + 4.0 = \mathbf{10.0}$$

The Decision: The agent chooses B.

The Agentic Logic: Even though A has a much higher predicted baseline score (8.0 vs 6.0), the agent realizes it knows almost nothing about the chemical space of B. Testing B provides maximum *Information Gain* to improve its neural network for the next round of synthesis!

AI Agents for Materials Discovery

Retrosynthesis and Materials

Science isn't just about medicine. We need novel materials: better batteries, biodegradable plastics, and advanced photoresists for chip manufacturing.

The Retrosynthesis Problem: The generative AI drew a picture of a perfect battery molecule. But how do we physically make it? What chemicals do we mix?

Agentic Planning: We treat Retrosynthesis as a Markov Decision Process (MDP). The agent starts at the target molecule (State) and takes "Actions" (applying chemical reaction rules) to step backward until it reaches cheap, commercially available starting materials.

IBM RoboRXN: AI + HPC + Robotics

IBM's RoboRXN [Pyzer-Knapp et al., 2022] demonstrates autonomous scientific discovery.

An AI agent uses a Transformer model to translate a target molecule into synthesis instructions.

These instructions are sent via the cloud to a robotic chemistry lab. The AI autonomously discovers, plans, and physically synthesizes a novel photoacid generator (PAG) without human intervention.

Self-Driving Labs

Architecting an Autonomous Lab

- We are now connecting all these modules into **Self-Driving Laboratories (SDLs)**.
- **Example Agent Loop:**
  - **Generative Module:** Proposes 10,000 new materials.
  - **Simulator Module:** Uses NNP/DFT to simulate properties, filtering down to 10 candidates.
  - **Robotics Module:** Mixes the chemicals for the top candidate.
  - **Analysis Module:** Physical sensors (mass spectrometers) measure the real-world result.
  - **Learning Module:** The real-world data is fed back into the Bayesian Optimizer to update the neural network, and the loop restarts.

Safety, Alignment, and Dual-Use

- If we give an autonomous agent the ability to design novel proteins and direct robotic labs to synthesize them, we face a massive cybersecurity and alignment risk.
- **Dual-Use Dilemma:** The exact same RL agent used to optimize a drug for low toxicity can have its reward function mathematically inverted to optimize for *maximum* toxicity (e.g., ricin).
- **The Engineering Imperative:** As AI engineers building these systems, you must design hard safety constraints (like the CMDPs we learned in Lecture 12) at the API level, ensuring the robotic lab physically rejects unsafe chemical synthesis requests.

Distributed AI Agents for Scientific Discovery (AI-BIOME)

AI-BIOME: An Autonomous Lab for Discovery of Bio-Inspired Materials

AI Agents are increasingly deployed in real-world robotics settings.

One of the most promising environments is automated laboratories to accelerate and open-source scientific discovery.

AIBIOME is a scientific hub that will automate experimentation for materials discovery

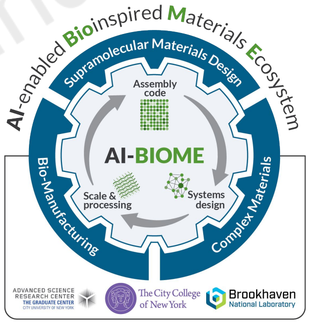

AI-BIOME Science Drivers & Cross-Cutting Capability

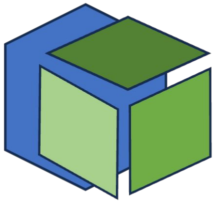

CC1: AI/computation/Lab Automation

SD 1: Assembly Code
SD 2: Complex Materials
SD 3: Scale & Process

Science-Driver: 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

SD1: Assembly Code

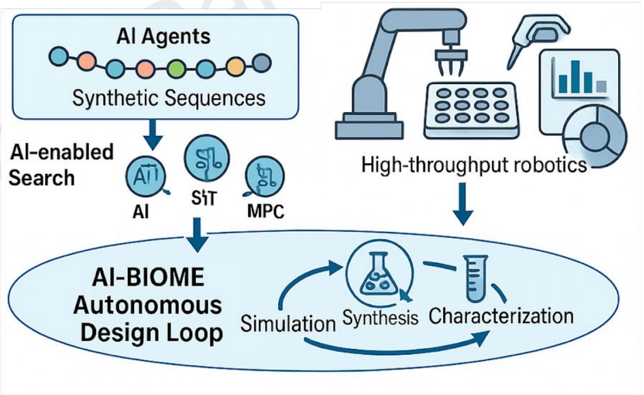

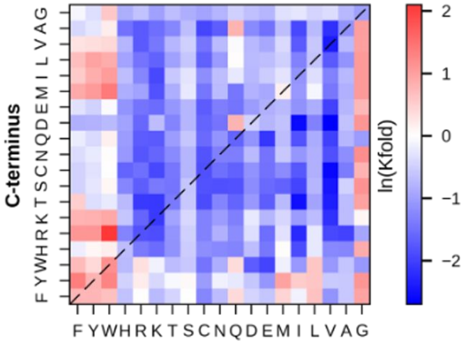

AI-BIOME

Science Driver: 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

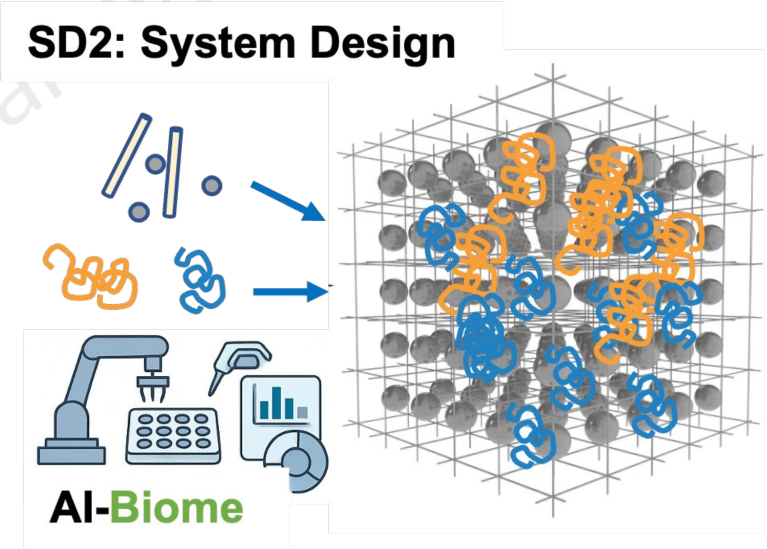

Science Driver: Distributed Scaling of Bio-Inspired Materials

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!

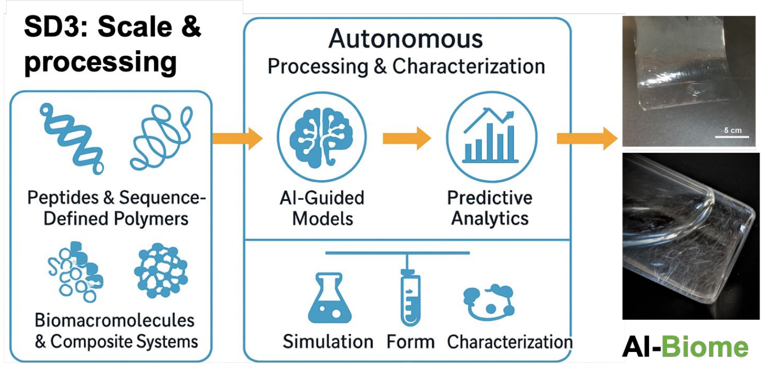

Questions?