A013 · Can AI Scientists Discover Better Drugs? Automating Objective Design, Property Prediction, and Molecular Optimization

Presenter: Yikun Zhang — Northeastern University

Authors: Yikun Zhang, Xiwei Cheng, Tianyu Liu, Yuanqi Du, Wengong Jin

Session: Day 2 · Fri Oct 2, 2026 · 2:15–4:15 PM

Can 'AI Scientists' discover better drugs, or only automate steps a chemist already knows how to take? Drug design, searching a chemical space of more than 10^{60} molecules, usually breaks into three stages: formulating the objective, predicting candidate properties, and optimizing molecules against it. Today each relies on humans to hand-build the underlying models, which is slow and often suboptimal. I present an AI Scientist that automates all three. On methicillin-resistant Staphylococcus aureus (MRSA), a WHO-designated urgent threat, it trained a property-prediction model that beat the Chemprop model our team hand-built for our 2023 Nature paper on the same data, then designed an antibiotic four times more potent than the human-designed lead, with a therapeutic index of 32. Two frameworks drive this: SAGA evolves the design objective through a bi-level loop of LLM agents, and DrugSAGE builds and refines state-of-the-art predictors and optimizers by reusing cross-task experience, surpassing the human leaderboard on Therapeutic Data Commons where prior agents cannot.

Can AI Scientists Discover Better Drugs? Automating Objective Design, Property Prediction, and Molecular Optimization

Yikun Zhang[1,2], on behalf of the SAGA and DrugSAGE teams
[1]Northeastern University [2]Broad Institute of MIT and Harvard
zhang.yikun@northeastern.edu

Motivation

Can "AI Scientists" discover better drugs, or only automate steps a chemist already knows how to take? Drug design searches a chemical space of more than 10^{60} molecules and usually breaks into three stages: formulating the objective, predicting candidate properties, and optimizing molecules against that objective. Today each stage relies on humans to hand-build the underlying models, which is slow and often suboptimal. We present an AI Scientist that automates all three, driven by two complementary frameworks, SAGA and DrugSAGE, and validated experimentally on a WHO-designated urgent pathogen.

SAGA: evolving the objective

Most discovery agents optimize a fixed, human-specified objective. For grand challenges that objective is only an imperfect proxy, and optimizers exploit the gap between the proxy and reality, scoring well while missing what an expert would catch. SAGA (Scientific Autonomous Goal-evolving Agent) treats objective design itself as the search problem. It runs a bi-level loop of LLM agents: an outer loop of a planner, implementer, and analyzer proposes new objectives and compiles them into executable scoring functions, while an inner loop optimizes candidate molecules under the current objective. Analyzing each round's outcomes lets SAGA revise the objective rather than treat it as a fixed input, and it supports co-pilot, semi-pilot, and autopilot modes so a scientist can steer as much or as little as they want.

DrugSAGE: self-evolving experience

Building state-of-the-art predictors and optimizers is an expensive search over tools, architectures, and training recipes. Existing agents can reach strong solutions through trial and error, but discard that experience and pay the full search cost on every new task. DrugSAGE (Self-evolving Agent Experience) keeps a cross-task memory of verified skills, statistical evidence about which strategies work, and a record of recurring errors and their fixes. The memory plugs into an MCTS-based search loop, with a formal guarantee that the memory-augmented selection policy preserves the regret bound of standard UCB. As memory grows the search narrows, and in some cases DrugSAGE transfers a working solution directly with no test-time search. Across 33 molecular-property-prediction tasks it ranks first among nine state-of-the-art agents; with memory from 16 smaller tasks it reaches an average normalized score of 0.935 on 17 held-out tasks and beats every baseline by 10–30% in a zero-test-time-search regime, surpassing the human leaderboard on Therapeutic Data Commons where prior agents cannot.

Wet-lab validation on MRSA

We tested the full pipeline on methicillin-resistant *Staphylococcus aureus* (MRSA), a WHO-designated urgent threat. DrugSAGE trained a property-prediction model that beat the Chemprop model our team hand-built for our 2023 Nature paper on the same data. SAGA then evolved the design objective and drove optimization to a candidate antibiotic that is four times more potent than the human-designed lead, with a therapeutic index of 32. Automating all three stages did not just reproduce the human workflow; it produced a better molecule than the expert baseline.

Takeaway

Formulating the objective, predicting properties, and optimizing molecules can all be automated, and doing so end-to-end can outperform the hand-built pipeline it replaces. Together SAGA and DrugSAGE point toward AI Scientists that improve the drug-design process itself, not only the steps within it.