BraintankAI

EXPLAINER · METACOGNITIVE RESEARCH AGENTS · 5 MIN

An AI research agent that keeps notes on what it knows before deciding what to try next

Shanghai AI Laboratory's Sage-Taiyi adds a supervisory layer that tracks findings, assumptions and open questions, then uses them to steer an AI's scientific idea search.

15 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Jordyn St John

Ask a capable language model for research ideas and it will produce plenty. Ask it to keep going, round after round, and a quieter problem appears. Each round mostly reacts to whatever happened last. An experiment that failed three rounds ago is forgotten. A contradiction between two lines of evidence is never resolved. Several promising ideas may all quietly depend on the same untested premise. Human scientists make these mistakes too, but they have habits for catching them: lab notebooks, group meetings, the nagging sense that something does not add up.

A 2026 paper from the Shanghai AI Laboratory, posted on arXiv as a preprint that has not yet been peer reviewed, tries to give an AI research system an equivalent habit. The system, called Sage-Taiyi, is built around metacognition, which means, loosely, thinking about one's own thinking. It keeps a running account of what its search has learned and lets that account decide where to look next and which reasoning methods to use.

The idea

The authors start from a family of systems that turned one-shot idea generation into loops. The AI Scientist runs through ideation, experiments and writing. Co-Scientist cycles through generating, critiquing and improving ideas. Robin coordinates several agents that reason with experimental feedback. The authors argue that iteration by itself does not guarantee that findings from different branches of a search stay attached to their evidence and the conditions under which they hold. Nor does it guarantee that those findings shape which methods the system picks later. When each round leans mainly on the latest feedback, they say, a system can repeat errors, keep digging down unproductive paths and over-react to recent results.

Their remedy borrows from cognitive science, specifically a model of metacognition proposed by the psychologists Nelson and Narens. In that model, mental activity is split into two levels. The object level does the work itself, here the science problem. The meta level watches that work and steers it. Two flows connect them. Monitoring updates the meta level's picture of how things are going. Control uses that picture to choose what to do next.

The authors also draw on older accounts of discovery as search through a problem space made of representations, operators and heuristics. They add work that treats choosing a strategy as a form of reasoning in its own right, sometimes called metareasoning. Put together, the pitch is that a research agent should reason not only about the problem but also about its own state of knowledge, and should choose its methods accordingly.

How this paper puts it to work

At the object level, Sage-Taiyi keeps three linked records. A Target Tree breaks the research question into smaller goals, each with local strategies and partial ideas. An Idea Graph works like a family tree of complete ideas, recording which were produced from which by five transformations the authors call Merge, Fuse, Diverge, Repair and Evolve. An Evidence State stores experimental and computational results, literature and expert input, each tied to its source.

The meta level holds two things. The first is a Cognitive State, a selective, source-linked set of running notes on findings, competing explanations, assumptions and unresolved questions. It also records judgements such as noticing that several ideas share one untested premise. The second is a reusable Cognitive Repertoire: ten research-strategy templates, such as analogical transfer, inverse design (working backwards from a desired property) and bottleneck-driven redesign. Alongside these sit smaller reasoning moves grouped under framing, constructing, evaluating and learning, tailored to each domain using classic papers in that field. Every cycle begins with Monitor, which updates the Cognitive State. When the next task calls for new ideas, Select picks heuristics and operators and Search applies them to a chosen goal or idea; otherwise the task goes to execution or validation. A budget and a stopping rule end the run.

A surrounding harness runs sub-agents in parallel and calls tools. One is Taiyi-Optimizer, the team's tool for Bayesian optimisation, a statistical method that chooses the next experiment to run based on results so far. Validation results flow back into the Evidence State.

What it shows, and what it does not

The main test is Taiyi-Bench, a benchmark the authors built themselves. It holds 265 tasks drawn from 50 papers across ten fields, from earth science and economics to materials science and communication engineering. The source papers were all first made public after the reported knowledge cutoff of GPT-5.6-Sol, the model Sage-Taiyi runs on, to reduce the chance it had already seen the answers. Each paper's key decisions are revealed step by step, and systems must propose ten ideas per task. Scoring covers how well a system rediscovers what the real study did and how novel, effective, feasible and varied its open-ended ideas are. An automated, literature-grounded grader applies scoring rubrics to the proposals.

Against ten alternatives, including plain prompting of three frontier models, the general coding agents Codex and Claude Code, and the scientific systems EvoScientist, AI Scientist-v2 and Robin, Sage-Taiyi scored highest overall: 53.95 against 47.85 for AI Scientist-v2. It came first on all eight measures and won about two-thirds of head-to-head comparisons with AI Scientist-v2. It also led on two outside benchmarks, AgentIdeaBench and IG-Arena. The authors also ran five collaborations with domain experts. In one, a cobalt-catalysed chemical reaction went from a trace of product to an 84 per cent yield after condition searching. Another reaction went from 13 to 69 per cent. A magnesium alloy was made and tested at a mean tensile strength of 342 megapascals. A metamaterial filter design outperformed an expert design, though only in simulation.

The limits matter. Taiyi-Bench and its grader come from the same team. The authors describe the grader's agreement with human experts as good but do not put a number on it. Scores are reported without error bars or significance tests. Novelty is the largest single source of the overall lead; on effectiveness, feasibility and identifying the reference study's root problem, the margins over the best competing system are under 1.5 points. In absolute terms, the system fully recovers the reference study's decisions in fewer than a fifth of cases. Comparison systems were cut down to their ideation stages, which may not reflect how they were meant to run. The leaderboard alone also cannot say whether the metacognitive layer, rather than more structure or more computation, drives the gains.

The laboratory results are single case studies. Experts chose which ideas to pursue and a Bayesian optimiser did the condition search. Apart from the simulated filter, no comparison shows what experts or the optimiser would have achieved alone. Two of the five collaborations were simulation only, and the authors note that mechanistic tests in the chemistry cases remain to be done. What the paper does show is a concrete design for an agent that keeps track of its own understanding, together with early, author-run evidence that it can feed productive research workflows. Later stories in this series will look at how far that idea stretches over longer investigations, and how to check the conclusions such runs reach.

PAPER	*Sage-Taiyi: A Metacognitive Architecture for Scientific Discovery*
AUTHORS	Weida Wang, Benteng Chen, Yigeng Jiang, Shuchen Pu, Weiping Shen, Zichao Liu, Yulin Su, Zitong Wang, Siyu Zhou, Zifu Wang, Yizhi Li, Xin Zong et al.
VENUE	arXiv preprint — not yet peer reviewed · 9 November 2026
LINK	arxiv.org/abs/2611.12762
INDEX	stacks.braintank.ai
CODE	github.com/Alpha-Innovator/Sage-Taiyi