BraintankAI ASIA

DEEP READ · MODEL-BASED SEARCH · 8 MIN

HONG KONG · Hong Kong University of Science and Technology

An LLM proposes designs and a statistical model picks which ones get tested

In a 2026 preprint, Large Discovery Models pair LLM suggestions with a statistical model that weighs predicted value against uncertainty to decide what earns a costly test.

9 November 2026 · arXiv preprint, not yet peer reviewed



Photo: National Cancer Institute

Zhongwei Yu, Jun Wang and colleagues call their system a Large Discovery Model (LDM), though it is less a single model than a loop with a division of labour. A large language model (LLM) proposes candidate designs, such as edits to a training program, antibody loop sequences or small molecules. A statistical model fitted to the results so far estimates how good each candidate is likely to be and how unsure that estimate is. Candidates that score well on that combined measure go for expensive evaluation. The work is a 2026 arXiv preprint, not yet peer reviewed, and the authors call it an initial v0.1 release.

The target is the shape of much scientific search. The spaces are vast and structured, and each test is costly, often noisy and opaque, so only a sliver of candidates can ever be tried. LLMs can generate valid candidates. The authors argue, however, that the models' likelihoods and self-assessments are poor stand-ins for what an experiment will measure and unreliable guides to

uncertainty, especially for unfamiliar designs. Inference-time reasoning methods succeed where cheap verifiers, such as a unit test, can grade attempts quickly and repeatedly. Discovery often lacks them. The paper's bet is that judgement should come from a model grounded in measurements, one that knows what it does not know.

How it works

The proposer LLM generates designs conditioned on a running context. That context can hold past results, the LLM's own reflections and the predictions of a surrogate, a cheap stand-in for the real test. The surrogate is a Gaussian process (GP), a statistical model that learns from a handful of tested designs. For any untested design, it gives a predicted score plus an uncertainty band that narrows as data accumulate. An acquisition function folds the two into one number saying how worth testing a candidate is. The authors use standard rules. Expected Improvement (EI) scores how much a candidate is expected to beat the current best. Upper Confidence Bound (UCB) adds a multiple of the uncertainty to the prediction. Expected Hypervolume Improvement handles cases where several objectives compete.

The authors start from an objective that rewards high acquisition scores but penalises straying far from the LLM's own tendencies. From it they show that the ideal search distribution is the LLM's proposal distribution, reweighted exponentially towards high-acquisition candidates. That cannot be sampled exactly. Instead, each round the LLM proposes a pool and the surrogate scores it. The system then either samples with probability rising with acquisition score or keeps the top scorer, a method known as best-of-N. Candidates arrive in one of two ways. In direct generation, the LLM writes complete designs and a validity checker filters them. In indirect parameterisation, the LLM picks which knobs to tune and their ranges, a separate sampler fills them in, and the GP models only that chosen subspace.

When progress stalls, a reflection step lets the LLM add or reshape search dimensions. The authors call this the discovery mechanism, and it anchors their theory. Each round's shortfall against the best possible design splits exactly into two parts. The first is a discovery gap: how far the best design within the LLM's current reach falls short of the true optimum. The second is an optimisation regret: how far the chosen design falls short of the best reachable one. In this picture, acquisition-guided selection works on the second term, while the LLM's proposals, which reflection can widen, govern the first. Optionally, the record of candidate pools and acquisition-weighted choices is used to fine-tune, or further train, the proposer. It then learns what the surrogate would value rather than raw scores. A reasoning-augmented variant first writes a surrogate-grounded rationale about progress before proposing.

What the evidence shows

All three tests are computational. The first is autoresearch, a public project released by Karpathy in 2026. In it, an AI coding agent edits a training script, and each edit is judged by training a small language model for five minutes. The score is validation bits per byte (val_bpb), a measure of how well the model predicts held-out text, where lower is better. LDM used indirect parameterisation and EI. It was compared with the original LLM-only loop under the same coding model, environment and five-minute-per-run budget on H100 graphics processors. Both started at 1.0069.

The baseline levelled off at about 0.9767 after roughly 255 runs, while LDM reached 0.9342, a reduction about 2.4 times as large. A separate LDM run on B200 hardware reached 0.9023 in 59 runs. That put it ahead of two autoresearch@home leaderboard entries, at 0.9264 and 0.9274. The edits LDM kept ranged from architectural changes to a narrower model and a different training optimiser.

Ablations, which vary one component at a time, pointed the same way. Scoring 64 candidate variants per round reached a better plateau than scoring 16. Freezing the searched features after the first iteration disables the discovery mechanism, and it left the search on a worse plateau. Choosing by predicted score alone was competitive early but levelled off at a worse value than either uncertainty-aware rule, with UCB best. A pure-LLM loop with no surrogate plateaued near 0.956 after 875 experiments.

The second test was antibody design. The target was CDRH3, a loop of about 11 amino acids that largely determines what an antibody binds. That gives some 200 trillion possible sequences. Each was scored by Absolut!, a simulated binding energy, against five antigens. Variants in which the LLM wrote sequences directly were only marginally better than standalone LLMs. The authors' policy-based variants did better and reached affinities comparable to AntBO, a method built specifically for antibodies. Random search, pure LLM generation and the combinatorial optimisers TURBO and COMBO were also compared. The third test optimised molecules against KRAS G12D, a mutated cancer-related protein. Two goals were pursued together: a docking score and a property the authors call neural activity. The score was hypervolume, a measure of how much of the objective space the best trade-offs dominate. Summarising its antibody and molecule results against LLM-only reflection or traditional statistical search, the abstract reports an 18.2 per cent relative decrease in binding energy. It also reports more than 60 per cent relative gains in multi-objective molecular performance.

Fine-tuning used Qwen3.5-9B as the proposer. On autoresearch over 80 iterations, the fine-tuned model with reasoning reached 0.9788. That was just ahead of the larger Qwen3-Coder-30B at 0.9801. The base model, without fine-tuning or reasoning, managed 0.9965. Fine-tuning also improved on the base proposer across the five antigens. The molecule task tests whether this carries over to a different problem. There, over 80 evaluations, the reasoning-augmented fine-tuned model reached the highest final hypervolume of five policies.

What it does not show

Nothing here has touched a laboratory. Binding energies and docking scores come from computational oracles, programs that stand in for real measurements. The authors list wet-lab validation as future work. Their definition of discovery covers only problems whose search space, evaluation and objective can be written down. A discovery also counts relative to the agent's prior knowledge, so it may be a rediscovery from the field's point of view.

The 2.4-fold figure compares LDM's endpoint with where the baseline stopped after about 255 runs. A separate pure-LLM loop ran 875 experiments and plateaued near 0.956. That narrows the gap, though the authors still describe it as well above LDM's 0.934. Searches like these are stochastic, so any margin is only as solid as the repeats behind it. The leaderboard entries come from outside

setups. The authors describe their run as hardware-matched, but the comparison is still suggestive rather than controlled. Table 1 credits the individual kept edits with improvements summing to about 0.084. It does not say which run they come from, so they cannot be checked against the 0.073 H100 reduction.

Other margins are modest. On antibodies, LDM matches AntBO rather than beating it. In fine-tuning, 0.9788 against 0.9801 is a thin lead over the larger model. The base model with reasoning (0.9825) also slightly beat the fine-tuned model without it (0.9829). On this benchmark, the reasoning step did about as much as the training.

Theory and practice also diverge. The regret analysis assumes, among other things, UCB acquisition, one evaluation at a time and a well-calibrated surrogate. It also places conditions on how closely the LLM's proposals cover the region around the optimum. The autoresearch runs, by contrast, used EI in batches with an added diversity term. There, too, the GP never sees programs. It sees only a lower-dimensional, LLM-chosen set of parameters, with its internal settings held fixed, which the authors adopt for stability when data are scarce. The surrogate's view of the problem is therefore bounded by whichever knobs the LLM exposes.

Where it fits

The paper sets itself against inference-time search that relies on cheap verifiers. That family runs from chain-of-thought, tree of thoughts and Self-Refine to AlphaZero-style search, AlphaEvolve, o1 and DeepSeek-R1. The authors cite autonomous research agents (Lu et al., 2024) and LLM-written research proposals (Si et al., 2024) as evidence that fluent generation is not enough. The proposals showed limited diversity and weaker feasibility, and LLMs did not reliably evaluate their own ideas. They also cite work by Gupta et al. and Chang et al. (2025) showing that LLM likelihood need not track measured properties.

Bayesian optimisation is a family of methods, such as EGO and GP-UCB, for choosing expensive experiments from few observations. From it, the authors take principled, uncertainty-aware selection. The LLM supplies what such methods lack on their own: a representation of the problem and a way to propose candidates. The authors also borrow the framing of infinite-armed bandits, where options cannot all be listed and must be drawn from a source. Here, the LLM is that adaptive reservoir. Among earlier hybrids of LLMs and optimisation, they describe OPRO as lacking calibrated uncertainty. They describe LLAMBO, BoChemian and LABO as earlier ways of building LLMs into Bayesian optimisation. Their claimed distinction is that acquisition value steers generation, compute and fine-tuning alike. Other approaches are steered by human preference, a verifier or a reward model's score.

As an opening to the question of how machines might discover, the paper isolates a sharp tension. Generators are expressive enough to propose almost anything, yet they cannot reliably say which proposals deserve an experiment. The paper's answer keeps the LLM as a source of options and hands the deciding vote to a model anchored in measurements. It also treats widening the search as part of the problem. Whether that holds outside simulation is a question the authors leave for future work.

PAPER	*Large Discovery Models: Empirically-grounded Model-Based Open-Ended Search*
AUTHORS	Zhongwei Yu, Yan Song, Xue Yan, Anjie Liu, Xingyu Lu, Yihang Chen, Huichi Zhou, Siyuan Guo, Luoyang Sun, Sihan Chen, Xiangning Yu, Jun Wang et al.
VENUE	arXiv preprint — not yet peer reviewed · 16 August 2026
LINK	arxiv.org/abs/2608.15669
INDEX	stacks.braintank.ai
CODE	github.com/yzailab/Large-Discovery-Models