BraintankAI

DEEP READ · EVOLUTIONARY SEARCH · 7 MIN

Training small language models on logged search steps from 355 optimisation tasks

A 2026 preprint fine-tunes small open models on the score-improving steps of AI-driven evolutionary searches, with average gains on test problems that are uneven and unreplicated.

11 November 2026 · arXiv preprint, not yet peer reviewed



Photo: National Cancer Institute

Some of the more striking AI results in mathematics and engineering have come from a simple loop. A program keeps a pool of candidate solutions, asks a large language model (LLM) to modify one, scores the result with an automatic checker, and repeats. Systems built this way have produced state-of-the-art results on maths conjectures, GPU kernels (small programs that run on graphics chips) and combinatorial puzzles; AlphaEvolve and the open-source OpenEvolve are two examples. In a 2026 preprint, not yet peer reviewed, Young-Jun Lee and Dongyeop Kang of the University of Minnesota and colleagues at Carnegie Mellon, KAIST and elsewhere ask whether the LLM's part in that loop can be taught directly to small models. Their method, Evolution Fine-Tuning (EFT), trains open models of 2 billion to 9 billion parameters on logged steps from searches on 355 of the 371 optimisation tasks the authors collected. Parameters are the internal values a model learns, and their number is a rough measure of its size.

The motivation is that each search starts from nothing and forgets what it learned. The model inside the loop is usually frozen. Such searches often rely on large proprietary models, because the loop needs good proposals at every step. The authors report that in their experiments, open models under 9 billion parameters struggled to follow these searches. Methods that do update a small model during a search throw the updates away when the task ends. The know-how of what to change, what to keep and when to backtrack therefore lives in the surrounding software, not in the model. EFT tries to move that know-how into the model's weights so it can be reused. The authors report average gains at every model size, but the gains are unevenly spread and no repeated runs are reported.

How it works

The outer program, called a scaffold, treats the LLM as its mutation operator, the component that proposes a changed version of an existing solution. The authors gathered 371 tasks from existing benchmarks. These include maths problems used by AlphaEvolve, competitive-programming and algorithm-tuning suites, GPU-kernel tasks, variants of Erdős problems and symbolic regression (finding a formula that fits data). Each task comes with a starting solution and a deterministic evaluator that scores any candidate.

They ran OpenEvolve with a far larger model, Qwen3.5-397B-A17B, as a teacher. It proposed both diff-based edits (patching part of a solution) and full rewrites, over several randomised runs per task. Every parent-to-child step was logged with three things. The first was what the model saw, including instructions, the parent, some search history and evaluator feedback. The second was what it wrote, and the third was how the child scored. The authors filtered out evaluator failures, steps that broke a working solution and over-long examples. That cut 172,997 records to 156,731, which the authors call the Finch Collection. Of these steps, 39.4% improved the score, 19.2% left it unchanged and 41.3% made it worse.

Training is supervised fine-tuning: the model is shown the situation the teacher faced and learns to reproduce the teacher's answer. The authors used only score-improving steps, drawn from 355 tasks, and held 16 tasks back. Symbolic-regression runs dominated the logs, so they capped how much each task contributed, which left 30,445 examples. They fine-tuned Qwen3.5 models at 2B, 4B and 9B, plus Qwen3 at 8B, for a single pass over the data. The resulting models are called Finch. An optional second stage applies KTO, a preference-learning method that pushes the model towards edits that raised the score and away from those that lowered it. The authors describe this as offline reinforcement learning, meaning learning from recorded rewards rather than live trial and error.

At test time Finch replaces the LLM inside OpenEvolve, with its weights frozen, for 100 iterations, and the best score is recorded. Finch can also serve as the starting point for test-time reinforcement learning on a single problem. For this the authors used nanodiscover, an open reproduction of a method called TTT-Discover that updates a model's weights while it searches.

What the evidence shows

Evaluation covered 22 tasks in five groups. Mathematical discovery includes the Erdős minimum-overlap problem, two autocorrelation inequalities, packing 26 circles and a Hadamard determinant problem. The other groups are algorithm engineering, systems performance, competitive programming scored from 0 to 100, and algorithmic heuristics. The main comparison pits each Finch model against its own base model in identical OpenEvolve runs. Averaged over the main results table, Finch beat its base by 1.56% at 2B, 3.40% at 4B, 3.17% at 8B and 10.24% at 9B.

Those averages hide a wide spread. Finch-9B beat its base by 65.09% on circle packing and by 21% on the Hadamard problem. Finch-4B gained 74.30% on a systems task called Transaction. Gains on the first autocorrelation inequality, however, ranged from 0.01% to 0.30%. On the Hadamard problem the 4B and 2B models got worse, by 61.96% and 16.22%. Competitive programming gives a clearer picture. The 9B average rose from 32.46 to 46.01, the 4B from 14.52 to 31.97 and the 2B from 0 to 6.12. The 8B barely moved, from 24.21 to 24.56, and slipped on three of the six problems.

The authors also compared Finch with proprietary models run in OpenEvolve, using scores taken from earlier papers. On the Erdős problem, where lower is better, Finch-9B reached 0.381100. Claude-Opus-4.6 scored 0.381880, and the best human result is 0.380927. Finch-4B beat Claude-Opus-4.6, Gemini-3-Pro and GPT-5 on Transaction. On circle packing, though, the best Finch scored 1.936 against Claude-Opus-4.6's 2.629.

The ablations, experiments that change one ingredient at a time, are more telling. Training Qwen3-8B on all retained steps, including those that left the score unchanged or made it worse, left it worse than the untrained base on all three maths tasks tested. On Erdős it scored 0.492, against 0.404 for the base and 0.381 when trained on improving steps only. Adding KTO lifted the competitive-programming average from 24.56 to 37.30 at 8B and from 31.97 to 36.30 at 4B. Training on 355 tasks rather than 15 improved held-out performance (on tasks excluded from training) by 14.1% on average across three tasks. In one case study, the authors report that Finch borrowed techniques from recommender systems and numerical optimisation, while the base model kept repeating one standard method.

What it does not show

No repeated runs or error bars are reported, so each figure appears to be the best score from one search. Many differences sit in the third or fourth decimal place, where, without repeated runs, signal cannot be told apart from noise. The base models do not even rank by size. On Erdős, the 2B base beat the 4B and 8B bases, and its 0.381737 beat the Claude-Opus-4.6 figure taken from earlier work. That sits awkwardly with the authors' premise that small open models are much weaker in this setting.

The separation between training and testing is loose. Only 16 of the 371 tasks were removed from training. The evaluation tasks are described as disjoint from the training tasks only as far as possible, and some come from the same benchmark suites as the training data. The table's 10.24% average for Finch-9B is calculated over 11 of the 22 tasks. The abstract presents a 10.22% average gain across all 22 tasks, and the main text gives 10.31% for Finch-9B. The proprietary scores were not rerun under identical settings.

The test-time learning result is the weakest. Under nanodiscover, Finch-8B and the untrained Qwen3-8B tied on both circle-packing tasks and differed by 0.000008 on Erdős. The paper's text claims a 3.2% Erdős gain that its own table does not show. The authors acknowledge that Finch-8B still trails the original TTT-Discover run with GPT-OSS-120B. They did not rerun the original, because it costs about 500 US dollars per task.

The authors note further limits. OpenEvolve was the only scaffold, used for both data collection and testing, so gains inside stronger scaffolds such as EvoX are untested. Test-time learning was tried only on maths, not on practical work such as kernel engineering. The method is text-only and writes each candidate in a single turn. Beyond those, Finch imitates a single teacher model, and for one of the two transfer case studies the authors only hypothesise why the behaviour arose.

Where it fits

The paper sorts earlier work into two camps. Search scaffolds such as AlphaEvolve, OpenEvolve, ShinkaEvolve, GEPA and EvoX improve solutions with a frozen model; AlphaEvolve is not open source. Learning-based methods, ThetaEvolve and TTT-Discover, update a model with reinforcement learning during a search, but one task at a time, and they discard the updates afterwards. EFT trains beforehand and across many tasks. It is meant to work with any scaffold or as a base for test-time learning. It borrows TTT-Discover's definition of a discovery: beating the previous best known solution within the iteration budget.

An earlier story in this series looked at the other side of the loop: pairing LLM suggestions with a statistical model that weighs predicted value against uncertainty to decide which designs earn a costly test. (earlier story) That approach sharpens the choosing around a fixed proposer. EFT tries to improve the proposer itself, and its dataset of 156,731 logged steps is a large part of the contribution.

On trust, the clearest evidence is the competitive-programming table and the ablation showing that training on all steps, including score-lowering ones, backfires. Gains in plain search hold on average but are uneven and unreplicated. The claim that EFT strengthens test-time learning is not borne out by the paper's own numbers, and transfer across fields rests on two case studies.

PAPER	*Evolution Fine-Tuning: Learning to Discover Across 371 Optimization Tasks*
AUTHORS	Young-Jun Lee, Seungone Kim, Minki Kang, Alistair Cheong Liang Chuen, Zerui Chen, Seungho Han, Taehee Jung, Dongyeop Kang et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 June 2026
LINK	arxiv.org/abs/2606.29082
INDEX	stacks.braintank.ai
CODE	github.com/Open-Galapagos/evolution-fine-tuning