BraintankAI

DISPATCH · EVOLUTIONARY SEARCH · 3 MIN

Evolving web searches alongside solutions lifts AI optimisation for two of three models

A 2026 preprint adds an evolving web-search loop to AI-driven evolutionary search, with large gains for one model, modest ones for another and losses for a third.

21 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Fulvio Ciccolo

AI systems that improve programs by trial and error can stall when progress depends on knowledge their underlying language model lacks. In a 2026 preprint, not yet peer reviewed, Young-Jun Lee of the University of Minnesota and colleagues describe EvoDuet, which evolves web-search queries alongside the solutions those queries are meant to inform. Across 21 optimisation tasks, the authors report gains for two of the three language models they tested, one large and one modest, and a loss for the third, Qwen3.5-9B.

The work builds on evolutionary search scaffolds: loops in which a large language model (LLM) proposes modified programs, an evaluator scores them, and the best become parents for the next round. Earlier in this series, logged steps from such searches were used to train small models (earlier story). The authors call these loops closed, because they draw only on the run's own history and the model's built-in knowledge. That makes them stall when progress needs outside

information, such as how to use a newer version of a software library. A search tool alone did not fix this. On a denoising task with GPT-5.6-Luna, in-loop search returned already-seen pages 88.1% of the time and stopped improving after iteration 50 of 100.

To test whether outside knowledge helps at all, the team first supplied task-relevant reference documents, called oracle documents, at every round on 31 tasks. Progress is measured as normalised discovery gain (NDG). This is the share of the gap between a starting solution and the best previously reported score that a run closes, so 100% means matching the prior best. Oracle documents raised average NDG by 10.3 percentage points for Qwen3.5-9B and by 4.0 for GPT-5.6-Luna. They also lowered it by more than a point on eight and four tasks respectively.

EvoDuet nests two loops: the outer one improves solutions, and the inner one improves the search queries that feed them. On each round, the model writes an account of what it knows and what remains open about the current solution. It then chooses between using no documents, reusing stored ones or running a fresh web search. A fresh search runs several rounds of queries aimed at those gaps. Rather than testing a solution built from every page, the model predicts the score each document would lead to and keeps the top-ranked ones, which keeps the inner loop cheap. With documents in hand, the outer loop can generate several candidates in parallel and keeps the best valid one.

The team added EvoDuet to an existing scaffold called OpenEvolve and ran it with one candidate per round. Overall NDG rose from 61.3% to 82.3% for Gemini-3.8-Flash and from 74.1% to 78.0% for GPT-5.6-Luna. For Qwen3.5-9B, NDG fell by 14.4 points, or by 4.7 with eight candidates per round. In 50 sampled revisions, that model left a retrieved method unused or implemented it wrongly 26% of the time. Pooled across models and settings, the largest median gain came on partially solved tasks: 17.5 points where the baseline had closed 20% to 40% of the gap, against zero where it had closed 80% or more. In GPT-5.6-Luna runs, the model's score predictions also tracked reality. Across 4,324 retrievals, predicted scores and the scores the resulting solutions actually earned had a rank correlation of 0.86, where 1 is a perfect match.

EvoDuet's best runs beat the previously reported best score on eight tasks and tied on three, at a cost of $11.93 to $161.33 per run. Several of those margins are tiny; on one task the score moved from 0.999970 to 0.999975. A clearer gain came on a task called Swap Reduction, where the count being minimised fell from 15,186 to 14,835, about 2.3%. On two tasks, the record programs built on published solutions that the search had retrieved partway through the run.

The main results are each the best of four runs rather than an average. The tests behind two key design choices, the retrieval gate and the two-loop coupling of search and solving, cover only two and three tasks respectively, both with one model. Retrieved ideas led to failed revisions in 29 of 82 analysed runs, and more than a third of fresh searches found nothing judged promising. Two questions remain open: whether open search can help models that struggle to use what they read, and whether the gains hold up when runs are averaged rather than picked at their best.

PAPER	*EvoDuet: Bilevel Co-Evolution of Web Searching and Task Solving for Scientific Discovery*
AUTHORS	Young-Jun Lee, Jinheon Baek, Soyeong Jeong, Minki Kang, Seungyeon Jwa, Jonghyun Choi, Seungho Han, Dongyeop Kang et al.
VENUE	arXiv preprint — not yet peer reviewed · 30 September 2026
LINK	arxiv.org/abs/2609.40340
INDEX	stacks.braintank.ai
CODE	github.com/Open-Galapagos/EvoDuet