

Gen-Searcher: An Image Generator That Googles Before It Draws

Kabui, Charles

2026-04-06

[Read at ToKnow.ai](#)

Gen-Searcher: An Image Generator That Googles Before It Draws

First search-augmented agent for knowledge-grounded image generation

- +16%** KnowGen benchmark gain over Qwen-Image baseline
- +15%** WISE benchmark gain Dual reward agentic RL
- 8B** Fully open-sourced model Code, data, and benchmark

April 6, 2026

ToKnow.ai

Researchers at CUHK, UCLA, and UC Berkeley released [Gen-Searcher](#), the first image generation agent trained to search the web before it draws. Current image generators only know what they learned during training. Ask for something recent or niche, and they hallucinate. Gen-Searcher fixes this with multi-hop web searches that collect both textual context and reference images before passing them to the generator. Built on [Qwen3-VL-8B](#), it's fine-tuned

on 10,000 curated search trajectories, then refined with agentic reinforcement learning using dual rewards: one measures whether the search found correct information, the other whether the generated image matches. On [KnowGen](#), a new benchmark spanning 20+ real-world categories requiring external knowledge, Gen-Searcher improves Qwen-Image by about 16%. On the WISE benchmark, the gain is about 15%.

The search agent transfers without retraining. Plug it into Seedream 4.5 and accuracy jumps about 16%; plug it into [Nano Banana Pro](#) and it gains about 3%. For anyone building with image generators, this means you can bolt on knowledge-grounded search without modifying the generator itself. News illustrators, educational content creators, and marketing teams working with current or obscure subjects no longer need to hope the model's training data covered their topic.

Generative models are shifting from closed knowledge systems to agents that look things up. Text-based search agents like [OpenSeeker](#) are already maturing. Gen-Searcher extends that pattern to images, with a full open-source release of code, data, models, and benchmark.

Sources: - [Gen-Searcher Paper \(arXiv\)](#) - [Gen-Searcher GitHub Repository](#) - [Gen-Searcher-8B Model \(HuggingFace\)](#) - [KnowGen Benchmark \(HuggingFace\)](#) - [Gen-Searcher Project Page](#)

***Disclaimer:** For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)*