MAJOR LABS · REPORT · JUNE 2026

The State of AEO

There is no such thing as "the AI answer." Ask ChatGPT, Claude, Perplexity, and Gemini the same question and they cite the same source only 16 percent of the time. Answer engine optimization is not one game. It is four. A firsthand read of how the four engines distribute their citations.

16%	**546**	**YouTube**
cross-engine agreement	distinct domains cited	most-cited source
310 vs 80	**1,398**	**~1.1%**
domains: Gemini vs OpenAI	citations analyzed	were Wikipedia

30 representative queries, four engines, every cited URL captured. Firsthand and open: github.com/major-matters/aeo-tracker.

Finding 1 — Four engines, four different answers

For each query we took the set of domains each engine cited and measured the overlap. Across 99 engine pairings over 30 queries, the mean overlap (Jaccard) was **0.16**. Pick two AI engines, ask the same question, and roughly one in six of the sources behind their answers is shared. The other five-sixths differ.

The answer you get depends nearly as much on which AI you ask as on what you ask.

Finding 2 — Every engine has a citation personality

The engines disagree not just on which sources but on how widely to look. Gemini drew on 310 distinct domains; OpenAI on 80, leaning on a short list of big tech-review publishers.

Engine	Cites	Domains	Character
OpenAI	232	80	narrowest; big tech-review media
Claude	331	142	niche, specialist sites
Perplexity	454	181	video-heavy (YouTube x54)
Gemini	381	310	by far the most diverse

Finding 3 — YouTube is the most-cited source

The single most-cited domain across the corpus was not a news site or Wikipedia. It was **youtube.com**, cited in 21 of 30 queries. Video, review media, and personal-finance explainers dominate the top of the table. Wikipedia was 1.1 percent of citations; forums and Reddit under 1 percent.

Finding 4 — A long tail, not a winner-take-all

The 30 queries pulled in 546 distinct domains, and the top 10 accounted for only 16.7 percent of citations. The citation pool is wide and shallow. Visibility is still winnable from outside the incumbent set, and the wide pool is exactly why the engines disagree so much: with no shared ranking, each model picks its own slice.

Method: 30 representative queries across product, how-to, technical, health, finance, and comparison categories, run against four engines via their citation-returning APIs; every cited URL captured. Read-only; a snapshot of a non-deterministic system, valuable across many queries and over time. Tool + query set: github.com/major-matters/aeo-tracker.

Major Labs builds open-source primitives and measurement for the agentic web. Cite: Major Labs (2026). The State of AEO. majorlabs.co/reports/state-of-aeo.