ON-PAGE.AI RESEARCH

# Do AI Assistants Cite Original Content? 793 AI Citations Measured Against the Pages That Rank

Eric Lancheres · On-Page.ai Research | July 15, 2026

ABSTRACT

A growing industry premise holds that AI search surfaces reward original content: contribute something new, the reasoning goes, and AI answers will cite you. We tested the premise directly. For the same 50 keywords measured in our first study, we captured every source cited by Google's AI Overviews (shown on 49 of 50 SERPs) and by ChatGPT with web search (150 answers, three per keyword), then scored 312 page-one organic results on the Information Gain Score, a 0–100 measure of how much a page adds beyond the pages ranking alongside it. Cited and uncited ranking pages were statistically indistinguishable on originality (medians 52 vs. 55.5; permutation p = .07, directionally lower for cited pages). AI Overviews included the SERP's most original page in 33% of eligible SERPs, below the 50% expected from random selection of the same citation counts. Three in four AI Overview citations pointed at pages not ranking on page one, led by YouTube (22% of all citation links) and Reddit; ChatGPT, which consulted the web in only 21% of runs, cited page-one results almost never (1% of its citations), and 85% of its cited URLs did not recur across three identical runs. Of the off-SERP pages the answers cited, 40% carried no scoreable article text at all (mostly video); the text pages among them scored modestly higher than the pages that rank (median 57 vs. 54). Citation by today's AI answer surfaces appears largely decoupled from the originality of what ranks, with one notable exception in legal content, where cited pages were markedly more original than uncited ones.

KEY FINDINGS

- Being cited by an AI answer is not a reward for originality. The ranking pages AI Overviews cited were no more original than the ones they skipped, and they included the single most original page on the SERP less often than random picking would.
- Three in four AI Overview citations point at pages that are not on the first page of Google's own results for the query they answer, and more of those links go to YouTube than to any publisher's website.
- ChatGPT answered four in five of these common search queries from its own knowledge, citing nothing at all. When it did search, its sources barely overlapped Google's results, or its own sources on a repeat run.
- Legal was the exception: there, AI Overviews really did cite the more original pages.

## 1. Background

Our first study [1] measured how much new information top-ranking pages add beyond the pages they rank alongside, and found the median top-3 page only moderately original. Since then, the practical question has shifted surfaces: AI Overviews now sit above the organic results for most informational queries, assistants answer the same queries directly, and an industry has formed around “AI visibility,” with original, first-hand content widely recommended as the way to earn citations in AI answers [2, 3].

That recommendation contains a measurable claim: if AI answer surfaces prefer original sources, the pages they cite should score higher on information gain than the pages they pass over. To our knowledge no published measurement has tested this against live AI answers. This study does, for two surfaces: Google's AI Overviews and ChatGPT with web search. It uses the same keyword set, metric, and scoring pipeline as the first edition, so the two studies can be read together.

## 2. Method

**Queries and capture.** The 50 keywords from the first study [1] (ten verticals, five keywords each) were submitted verbatim on July 15, 2026 (US, English) to two surfaces. For Google, we recorded the SERP: the AI Overview with every source it cites, plus the organic results shown on page one, through an established SERP data provider [4]. For ChatGPT, we asked each keyword three times through the OpenAI API with web search enabled [5], recording the sources the answer actually cites inline. Whether the model chose to search at all was left to the model, as in the product, and recorded. “Cited” throughout means a source linked from the answer itself, not pages merely consulted during search.

**Scoring.** Every page-one organic result was scored with the Information Gain Score [1]: a 0–100 measure of how much of a page's main content adds beyond the pooled content of the pages currently ranking for the same keyword, compared by meaning rather than exact wording. Of 370 distinct ranking pages, 312 yielded scoreable main content; 58 (16%) did not, chiefly forum threads, video pages, and other thin-text formats, an exclusion that recurs below as a finding in its own right. Pages the assistants cited from outside the rankings were scored the same way against the same keyword cohorts (§3.7; the deepest citation lists were sampled to six per keyword, 281 pages in all).

**Analysis.** The primary comparison is between ranking pages the AI Overview cited and ranking pages it did not. This comparison stays inside each SERP, so keyword difficulty and vertical drop out. We report score distributions, a median difference with a permutation test that shuffles citation labels within each keyword, and per-SERP paired deltas. ChatGPT's citations are reported descriptively: they overlapped the rankings too rarely (three pages in the whole sample) for the same within-SERP comparison.

## 3. Results

### 3.1 How often AI answers appear and cite

Google showed an AI Overview on 49 of the 50 SERPs (98%), citing a median of 11 sources each, for 584 cited links in all. The single holdout was a marketing keyword. Coverage this close to universal, on a keyword set built for breadth rather than for triggering AI features, is itself

worth recording: for these informational and commercial queries, the AI answer is now the default page-one experience, and the organic results beneath it averaged eight visible slots.

ChatGPT behaved differently. Given the same queries, it consulted the web in 32 of 150 runs (21%) and answered the rest from its own knowledge, with no citations at all. When it did search, it cited 7.8 sources per answer on average (213 distinct URLs across the study). Four in five of these common search queries produced a ChatGPT answer citing nothing; for those, there is no citation to win.

### 3.2 Where citations point

Of the AI Overview's 584 citation links, 11% pointed at one of the keyword's top-3 organic results and 25% at any page-one organic result. The remaining 75% pointed at pages not ranking on the first page of Google's own results for that query (a further 16 links shared a domain with a page-one result without being that page). One likely mechanism deserves naming: Google has described its AI search experiences as issuing many related sub-queries behind a single question, known as “query fan-out” [6]. A page cited from beyond this query's page one may rank well for a related sub-query we never see. Our measurement is anchored to the visible query's SERP: the one a searcher would compare the AI answer against.

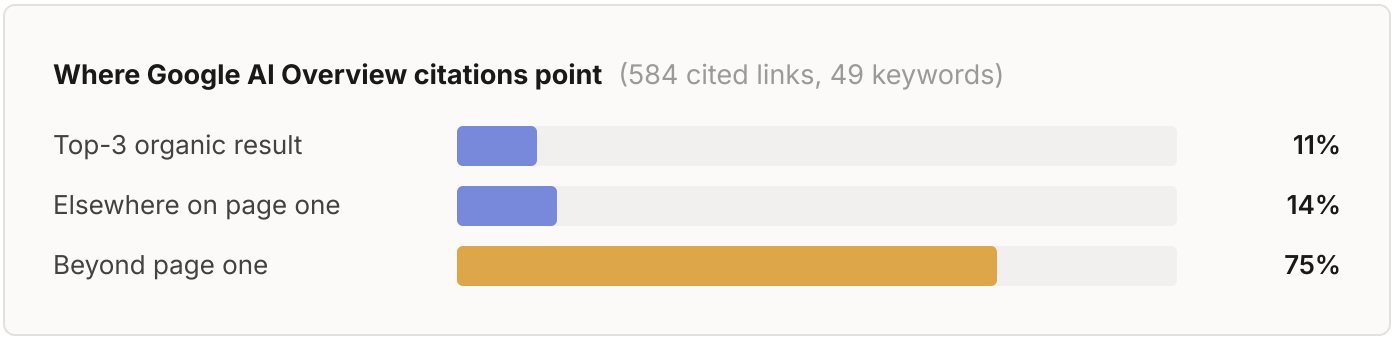


Figure 1. Share of AI Overview citation links pointing at the keyword's own top-3 organic results, elsewhere on its first page of organic results, or at pages not ranking on page one at all. Shares are of all 584 cited links across the 49 keywords that showed an AI Overview.

The destinations concentrate in a familiar direction: YouTube received 127 of the 584 links, or 22% of everything the AI Overviews cited, and more than four times the next domain. Reddit followed with 27. Facebook and Instagram together took 12. The most-cited destinations of

Google's AI answers, on this sample, are predominantly video and community platforms. Those formats also dominate the 16% of ranking pages our pipeline could not score for lack of article text.

| AI Overviews | Links | ChatGPT (searching runs) | Links |
|---|---|---|---|
| youtube.com | 127 | support.google.com | 9 |
| reddit.com | 27 | squaremouth.com | 8 |
| facebook.com | 7 | breville.com | 6 |
| instagram.com | 5 | law.cornell.edu | 6 |
| runrepeat.com | 4 | woocommerce.com | 5 |
| geeksforgeeks.org | 4 | forbes.com | 4 |

Table 1. Most-cited domains by citation links received. Counts are citation links, not originality scores; no third-party page is scored publicly in this study.

ChatGPT's citations point somewhere else entirely: only 1% of its 213 cited URLs were page-one organic results on Google for the query, representing three pages in the entire study. One caveat cuts this comparison some slack: ChatGPT retrieves through its own search stack, which draws on sources beyond Google's ranked results and only sometimes surfaces the same pages Google ranks. Low overlap with Google's rankings is not, by itself, surprising. The practical point survives the caveat: ranking on Google does not, by itself, put a page in ChatGPT's answers. Its list is led by official documentation and product pages rather than publishers, and its long tail is long: no domain received more than nine links.

### 3.3 Cited vs. uncited ranking pages

The premise under test says the cited rows of Table 2 should sit well above the uncited rows. They do not.

| Ranking pages | n | Mean | P25 | Median | P75 |
|---|---|---|---|---|---|
| **Cited by the AI Overview** | **146** | **52.3** | **40** | **52** | **65** |
| Not cited | 166 | 53.1 | 40 | 55.5 | 69 |

Table 2. Information Gain Score distributions for page-one organic results the AI Overview cited vs. did not cite. Median difference −3.5 points (cited lower); two-sided permutation test shuffling citation labels within keywords, p = .07.

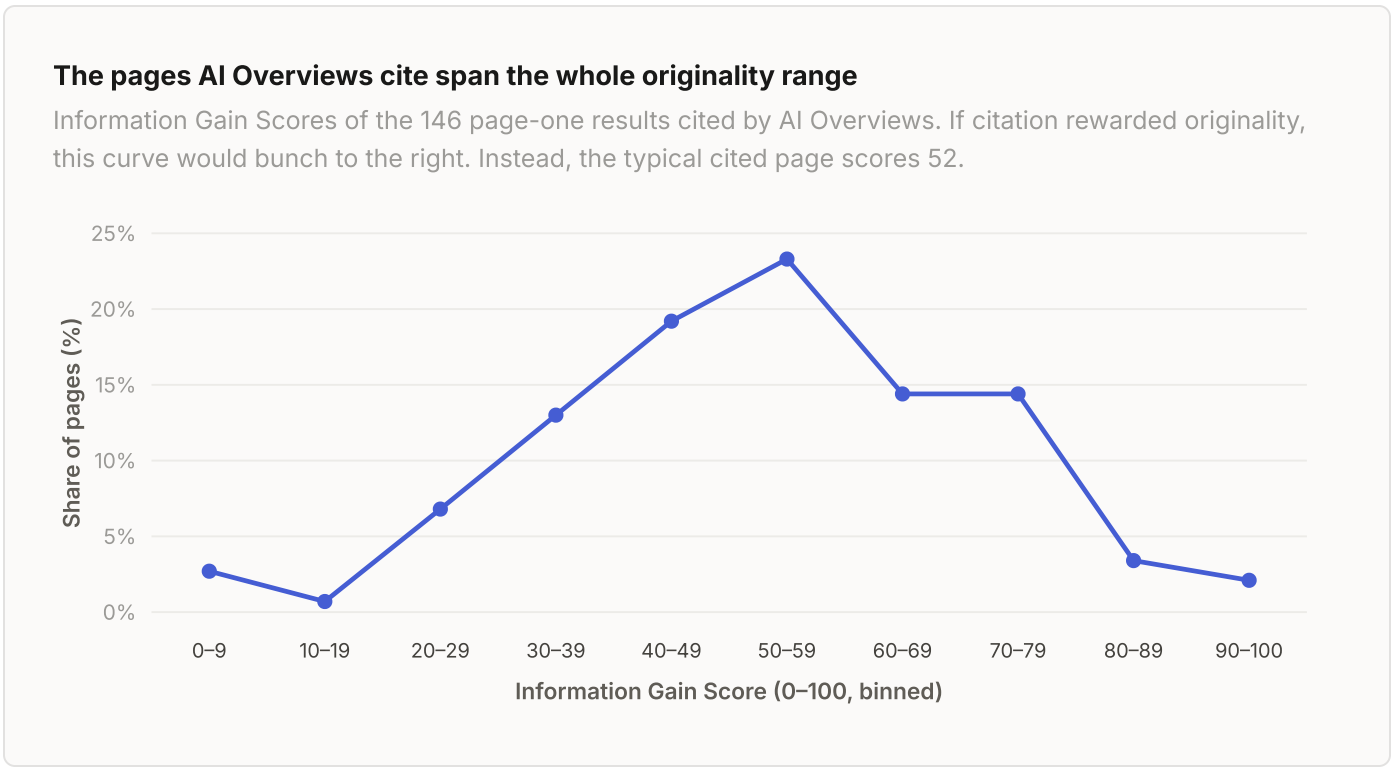


Figure 2. Information Gain Score distribution of the 146 AI-Overview-cited ranking pages, as the share of the group per 10-point bin. Cited pages span the full range, with the typical page at 52, statistically indistinguishable from the pages the Overview skipped (Table 2).

Cited pages were, if anything, slightly less original (median 52 vs. 55.5), a difference that does not reach significance and whose direction is the opposite of the premise. The within-SERP picture matches: in the 44 SERPs containing both cited and uncited scored pages, the cited side averaged higher originality in 19, essentially a coin flip. Grade composition tells the same story: 20% of cited pages graded highly original against 24% of uncited pages, and the share of mostly-shared pages was near-identical (23% vs. 25%). Unique numeric data points, the strongest correlate of the score in the first study, showed medians of 4 (cited) against 5 (uncited).

In short: knowing a ranking page's Information Gain Score tells you essentially nothing about whether the AI Overview cited it.

### 3.4 Does the answer cite the SERP's most original page?

A weaker version of the premise would settle for the AI answer reliably including the single most original result. It does not do that either. In the 46 SERPs where the AI Overview cited at least one scored ranking page, the highest-scoring page was among its citations 15 times (33%). Because the Overviews cited a median of three scored ranking pages per SERP, random selection of the same counts would include the most original page about 50% of the time. The observed rate sits below chance, not above it. The median originality rank of a cited page within its SERP was 4th.

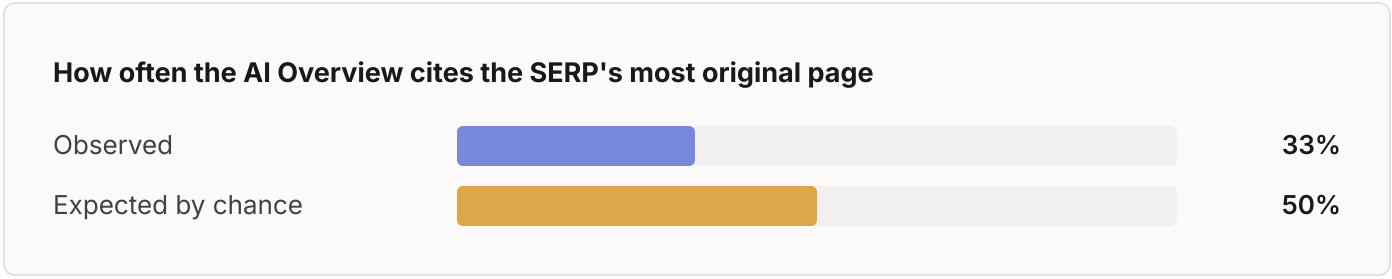


Figure 3. Of the 46 SERPs where the AI Overview cited at least one scored ranking page, the page with the highest Information Gain Score was among the cited pages in 15 (33%). Citing the same number of pages per SERP at random would include the most original page 50% of the time.

### 3.5 Two assistants, two different webs

For the 15 keywords where both surfaces produced citations, the overlap between what ChatGPT cited and what the AI Overview cited was effectively zero: mean Jaccard similarity of 0.0 at the URL level and 0.05 at the domain level. The two systems answering the same query cite different sources from what ranks, and from each other.

ChatGPT does not even reliably agree with itself. Across three identical runs per keyword, 85% of its cited URLs appeared in only one run; 3% appeared in all three. A page cited by ChatGPT today is, on this evidence, unlikely to be cited for the same query tomorrow. That complicates both the promise of “winning” a ChatGPT citation and any tool that reports one as a stable asset.

## 3.6 Variation by vertical

| Vertical | Cited median | n | Uncited median | n |
|---|---|---|---|---|
| **Legal** | **62** | **26** | **49** | **9** |
| Travel | 56.5 | 16 | 56 | 9 |
| Food & recipes | 58 | 10 | 65 | 17 |
| Home improvement | 64.5 | 8 | 72 | 12 |
| Marketing / SEO | 52 | 13 | 58 | 21 |
| Technology / dev | 51 | 17 | 57 | 14 |
| B2B SaaS | 47 | 18 | 53 | 21 |
| Personal finance | 44 | 14 | 47 | 23 |
| Ecommerce | 44 | 9 | 62 | 19 |
| Health | 40 | 15 | 53 | 21 |

Table 3. Median Information Gain Score of AI-Overview-cited vs. uncited ranking pages by vertical. Cells are small (8–26 pages); read as direction, not as precise estimates.

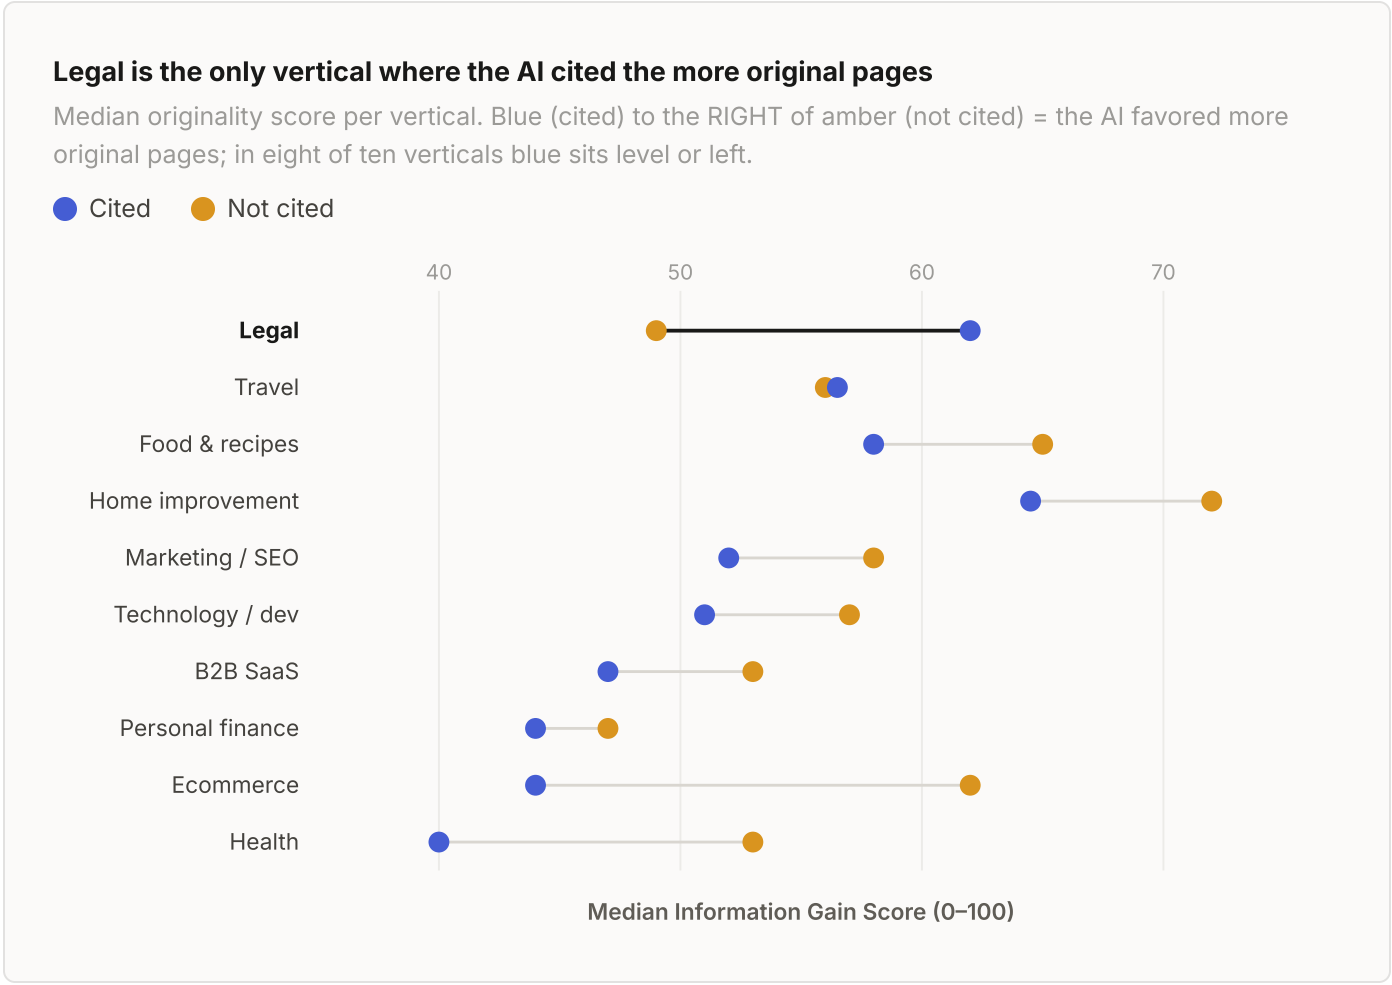


Figure 4. Median Information Gain Score of AI-Overview-cited (blue) vs. uncited (amber) ranking pages by vertical, on a shared 35–75 scale. In eight of ten verticals the cited dot sits at or left of the uncited dot. Legal (bold) is the reversal: cited pages markedly more original. Cells are small (8–26 pages); read as direction.

The null result is not uniform. In eight of ten verticals the cited median sits at or below the uncited median, with ecommerce most starkly (44 vs. 62). Legal runs the other way: pages the AI Overview cited graded markedly more original (62 vs. 49), and legal is also where the Overview cited the most ranking pages (26 of 35 scored). One vertical of five keywords is a direction, not a law; but it is consistent with higher-stakes topics drawing citations toward substantive sources, and it is the one place in this sample where the industry premise held.

### 3.7 The pages AI cites from beyond the rankings

Where do the other three-quarters of citations lead? We scored the sampled off-SERP cited pages (281; up to six per keyword) against the same keyword cohorts. The first result is about format: 113 of them (40%) could not be scored at all for lack of article text, and 98 of those were YouTube pages (the remainder mostly social profiles and forum threads our crawler could not

extract). For comparison, 16% of ranking pages were unscoreable. A large share of what AI answers cite is simply not article content, so its information gain, in the sense this metric measures, is undefined rather than low.

| Pages | n | Mean | P25 | Median | P75 |
|---|---|---|---|---|---|
| Ranking pages (all scored) | 312 | 52.7 | 40 | 54 | 67.2 |
| **Off-SERP cited pages (scored)** | **168** | **55.5** | **41** | **57** | **70.2** |

Table 4. Information Gain Score distributions: all scored ranking pages vs. scored off-SERP cited pages (text pages only; the 40% unscoreable share above is excluded by construction). Off-SERP sample capped at six per keyword.

The text pages AI answers do cite from beyond the rankings are modestly more original than the pages that rank (median 57 vs. 54), and a larger share of them graded highly original (28% vs. 22% of ranking pages). The small set of off-SERP articles cited by ChatGPT stands out further, with a median of 66 across its 14 scoreable pages, though the sample is too small to lean on. Read together with §3.3, the picture is consistent: within the rankings, citation ignores originality; but when an answer engine reaches beyond the rankings for citable text, what it retrieves skews slightly more original than page one, when it is retrieving article text at all.

## 4. Discussion

Two readings of these results should be kept apart. The first is about mechanism: nothing here shows that AI answer surfaces penalize originality, or measure it and discard it. The citation patterns are consistent with selection on other properties entirely: format (video, community discussion), source-type priors, or the breadth of retrieval behind a single answer, where many related sub-queries reach far beyond the visible SERP [6], with originality simply not among the operative criteria. A null is precisely what one would expect if citation selection never looks at information gain.

The second reading is practical, and it is the one the industry premise runs into. “Add original information and AI answers will cite you” is, on this evidence, not how today's surfaces behave for everyday informational queries: the Overviews cited less-original ranking pages as readily as more-original ones, skipped the most original page on the SERP more often than chance, and spent three-quarters of their links beyond page one, dominated by platforms rather than

publishers. Originality remains associated with ranking well (the first study's position gradient [1]), and ranking on page one is still where a quarter of Overview citations land. But as a direct lever on being cited, information gain currently has no measurable purchase, except, intriguingly, in legal.

The instability finding cuts at a different target. A citation that appears in one ChatGPT run in three is not an asset a dashboard can meaningfully track week to week; teams evaluating AI-visibility tooling should ask how a reported “citation” was sampled, how often it recurs, and across how many runs.

## 5. Limitations

All captures and scans were taken on a single day, in one locale (US, English), on one keyword set of 50 hand-built keywords chosen for breadth, not randomly sampled; AI surfaces change quickly, and these are July 2026 snapshots. ChatGPT was queried through its API with web search enabled, which may differ from the consumer product in retrieval behavior; three runs per keyword bounds, but does not eliminate, sampling noise in an answer surface this unstable. The AI Overview data reflects what the SERP showed our provider at capture time; page-one organic slots averaged eight, so “beyond page one” can include pages ranking just below the visible results, or ranking well for related sub-queries the answer engine issued behind the scenes (“query fan-out” [6]), which this study does not observe. Citation is measured as linking, not influence: a source can shape an answer without being linked. Unscoreable formats (video, forums) could not receive Information Gain Scores, so the originality of 12% of ranking pages, and a larger share of cited destinations, is unmeasured, a selection effect that this study reports rather than escapes. The Information Gain Score is our own metric [1]; it measures semantic contribution relative to a ranking cohort, not any search or answer engine's internal signal.

## 6. SEO Implications

For teams allocating content effort while AI surfaces mature:

- Keep the originality case honest: on current evidence, information gain is a ranking-side and reader-side investment, not a direct lever on AI citations. Pages that rank page-one still capture a quarter of Overview citations. Reaching page one remains the most measurable path into the

answer box.

- Treat formats as a separate axis: the largest single destination of AI Overview citations was YouTube, and community threads follow. Where a query's answer surface leans on video and discussion, a text page, however original, is competing in the wrong format.
- Audit AI-visibility claims with recurrence in mind: most ChatGPT citations in this sample did not survive a repeat of the same query. A citation report without a stability measure is a snapshot of noise.
- Watch vertical differences rather than averages: in legal, cited pages really were more original. If your vertical behaves like that, the originality-to-citation path may be live for you even though it is absent in the aggregate.

## References


1. Lancheres, E. (2026). “Information Gain in Google's Top-Ranking Results: A Measurement of 150 Pages.” On-Page.ai Research. api.on-page.ai/research/information-gain-study
2. “Creating helpful, reliable, people-first content,” Google Search Central documentation. developers.google.com/search/docs/fundamentals/creating-helpful-content
3. “AI features and your website,” Google Search Central documentation. developers.google.com/search/docs/appearance/ai-features
4. Google AI Overview results captured via SerpApi. serpapi.com/ai-overview
5. “Web search,” OpenAI API documentation. developers.openai.com/api/docs/guides/tools-web-search
6. “AI Mode in Google Search: Updates from Google I/O 2025,” Google (describes the “query fan-out” technique). blog.google/products-and-platforms/products/search/google-search-ai-mode-update



**Cite this study:** Lancheres, E. (2026). Do AI Assistants Cite Original Content? 793 AI Citations Measured Against the Pages That Rank. On-Page.ai Research. https://api.on-page.ai/research/ai-citation-study