



Generative **Pulse**

| **MUCK RACK**

What is AI Reading?

October 2026

Executive Summary

- **Earned media remains central to GEO:** About **84%** of citations come from earned media, while roughly **99%** come from non-paid sources. Journalism still represents about **22%** of citations, even as its overall share has declined.
- **No single source strategy reaches every AI provider:** Only **212 domains** appeared across all four providers. Communicators need a portfolio that includes owned content, earned coverage, national and specialist outlets, and distribution channels the relevant providers can access.
- **The question determines which content gets cited:** Practical guides account for **68%** of citations on how-to questions, while best-of prompts favor buyer's guides and reviews pages at **45%**. Industry-trend questions favor expert analysis, and comparison questions rely most on evergreen reference content. GEO strategies should align content formats with the questions audiences ask.
- **Journalism and owned media serve different purposes:** Nearly half of cited journalism covers breaking or recent news, while another **30%** provides reviews and buyer guidance. Owned-media citations are led by product and service pages at **58%**, followed by practical guides at **24%**. Effective GEO combines third-party validation with useful information published by the brand.
- **Reliable GEO measurement requires scale:** When providers answered the same prompt twice, source overlap ranged from **41%** to **94%**. Teams should use stable prompt sets, run questions repeatedly, report results by provider and track ranges over time rather than treating one response as definitive.

Methodology

We analyzed more than 40 million AI responses

This study examines the sources major AI providers cite in response to realistic prompts, including citation frequency, page type and the balance between earned and owned media.

We combined two data sources. Dedicated research collections ran primarily unbranded consumer prompts across ChatGPT, Claude, Gemini and Perplexity at set points in 2026. Muck Rack's Generative Pulse continuously ran mostly brand-focused prompts to track sourcing changes over time. Together, the datasets include more than 40 million responses across industries and question types.

We classified cited URLs by source and content type. Methods and coverage varied by dataset and are noted on the relevant charts.

Models included GPT-5, GPT-5 nano, Gemini 2.0 Flash, Gemini 2.5 Flash, Claude Haiku 4.5 and Perplexity Sonar Pro.

Providers introduced and retired models during the study, so comparisons over time reflect product and model changes as well as elapsed time. The findings describe observed behavior and may shift as providers update their systems.


We assigned cited links into categories as follows:

- **Journalistic:** News sites, and other journalistic coverage
- **Corporate Blogs and Content:** Third party corporate blogs and content **not** owned by a company/product targeted in the query
- **Owned Media:** Corporate Content created by a company/product targeted in the query
- **Press Release:** Press release published on any site
- **Academic/Research:** Scientific journals, arXiv, patents, research papers
- **Government/NGO:** .gov sites, public agencies, non-profit organizations
- **Paid/Advertorial:** Sponsored content, pages with marketing intent or sales landing pages
- **Social/UGC:** Social platforms (e.g., LinkedIn, Reddit) and other sources of user-generated content
- **Aggregators/Encyclopedic:** Sites like Wikipedia, Visual Capitalist and Britannica.com

**Which types of
media does AI cite?**

Despite reshuffling between categories, about 99% of links cited by AI come from non-paid media


Only around 1% of all links cited by AI came from paid sources like press releases and paid/advertorial content. These figures are largely unchanged since our May 2026 report.



84% of citations come from earned media


Our definition of earned media includes journalism, academic research, government and NGO sources, encyclopedic sites like Wikipedia, social platforms, and third-party corporate content.

Owned and paid content together account for about **17%** of citations, an increase of about **1 percentage point** compared to the May 2026 report.



Journalistic sources drops to about 22% of total AI citations

Journalistic sources dropped from about **27%** of all AI citations in May 2026 to **22%** today. These **7 percentage points** were mostly redistributed among other earned media types, which is why the percentage of earned media is mostly unchanged.




The four main providers are barely reading the same internet

ChatGPT draws 35% of its citations from government and NGO sources, while owned media dominates Gemini at 34%. Claude relies most on third-party corporate content at 24%. And Perplexity cites journalism the most often at 28%.

What each provider actually read in September 2026

Share of classified citations, 1,020 prompts per provider



Press release: 2% (ChatGPT), 0.2% (Claude), 0.1% (Gemini), 0.2% (Perplexity)
Paid/advertorial: 0% (ChatGPT), 0.2% (Claude), 0% (Gemini), 0.1% (Perplexity)

Tech platforms: 0.7% (ChatGPT), 0.3% (Claude), 0.2% (Gemini), 0.2% (Perplexity)
Social/UGC: 0.3% (ChatGPT), 3.3% (Claude), 5.1% (Gemini), 4.3% (Perplexity)


The same questions produced a very different source mix five months later

Source preferences shifted differently across providers. ChatGPT moved toward government and NGO sources, rising from **10%** to **35%**, and owned media, rising from **14%** to **20%**. Perplexity's reliance on third-party corporate content fell from **24%** to **20%**, while Gemini's increased from **14%** to **17%**.

Note: ChatGPT and Gemini used different model versions in April and September, so the comparison reflects model updates as well as elapsed time.

The same 102 questions, five months apart

Source mix on the longitudinal panel: April 2026 (WIAIR3) vs September 2026 (fresh run)



**What determines if
content gets cited?**

Citation behavior depends as much on the prompt as the provider

Perplexity cited sources in **100%** of answers, while Gemini remained above **92%** across every question type. ChatGPT's citation rate ranged from just **27%** for comparison questions to **96%** for industry-trend prompts, and Claude also varied widely by question.

This means the mix of prompts can strongly influence any citation-based measure of AI visibility.

How often each provider cites anything at all

Share of answers containing at least one citation, by question type

	ChatGPT	Claude	Gemini	Perplexity
Industry trend	96%	97%	99%	100%
Best of	75%	72%	92%	100%
Risk due diligence	38%	52%	98%	100%
Problem discovery	46%	44%	94%	100%
Comparative evaluation	27%	53%	98%	100%
How to	39%	34%	92%	100%

Each provider favors a different type of page

Product and service pages account for **50%** of Gemini's citations. ChatGPT stands out for product/service pages at **21%**, and data and research at **16%**, while Claude favors buyer's guides and expert analysis at about **20%** each. Perplexity leans most heavily on how-to content at **26%**. Despite these differences, breaking news represents just **2%** to **5%** of citations across providers.

What type of page gets cited

Share of each providers' citations by content type

	ChatGPT	Claude	Gemini	Perplexity
Product/Service page	21%	6%	50%	7%
Practical guide/How to	12%	17%	11%	26%
Buyer's guide/Review	6%	21%	8%	18%
Expert analysis/ Thought leadership	11%	20%	8%	14%
Evergreen reference	12%	15%	8%	14%
Data/Research/ Academic	16%	12%	2%	8%
Regulatory/Policy/Legal	11%	2%	3%	2%
Breaking/Recent news	5%	5%	2%	3%
Community/Discussion	0%	1%	8%	6%
Press release/ Official announcement	5%	1%	0%	0%

The prompt decides the type of content cited

AI does not favor one universal content type. Practical guides account for **68%** of citations when the prompt is a how-to question, while expert analysis leads industry-trend prompts at **41%**. Best-of questions favor buyer's guides at **45%**, and comparative questions lean most on evergreen reference content at **37%**. The strongest content format depends on the question a user is asking.

Share of citations by content type within each query type


Problem discovery	28%	16%	6%	15%	23%	4%	2%	2%	3%	0%
Best-of/ Recs	8%	45%	2%	10%	26%	2%	1%	1%	2%	0%
Industry trend	2%	4%	41%	9%	7%	18%	14%	1%	2%	2%
How-to	68%	3%	1%	5%	16%	1%	1%	4%	1%	0%
Comparative	6%	21%	11%	37%	14%	1%	4%	1%	4%	0%
Risk/Due diligence	36%	9%	14%	17%	9%	5%	4%	4%	3%	0%
	Practical guide/ How-to	Buyer's guide/ Review	Expert analysis/ Thought leadership	Evergreen reference	Product/Service page	Breaking/Recent news	Data/Research/ Academic	Regulatory/Policy/ Legal	Community/ Discussion	Press release/ Official announcement

Each provider pulls from a different long tail of sources

The models draw from broad but substantially different source pools. Across **1,020 prompts** per provider, the number of distinct cited domains ranged from **2,059** for ChatGPT to **11,543** for Perplexity. Only **212 domains** appeared across all four providers, showing that there is no single source ecosystem shared by every provider.

The universal set of citations is led by forbes.com, linkedin.com, energy.gov, epa.gov, fda.gov.

Distinct domains cited across 1020 prompts




The same prompt often produces a different mix of sources


When asked the same question twice, Perplexity returned nearly the same sources in both runs, with **94%** overlap. That fell to **56%** for Claude, **45%** for Gemini and **41%** for ChatGPT, meaning one run of a sample prompt isn't enough. Even when the specific sources changed, each provider cited sources at about the same rate and volume.

When it comes to AI visibility, many prompts and repeated runs, not a single answer, are the keys to success.


Overlap of cited domains between two runs of the same prompt on the same day




ChatGPT



Claude



Gemini




Perplexity

Insights & trends

ChatGPT leads in deprioritizing news

According to data from [Generative Pulse](#), ChatGPT is a major driver in the shift away from journalistic citations. While Claude and Gemini had mostly steady journalism citations patterns over most of the year, ChatGPT had two major adjustments. First, a big drop in April, and a second, more dramatic drop in July.

Weekly share of journalism citations, by provider




Major publishers disappeared from ChatGPT in a series of sudden waves

Reuters fell to nearly zero in January, followed by the Financial Times in March, several publishers in May and AP and Axios in July. Some of these publishers later returned to results. During the same period, citations to some other outlets—like, Harvard Business Review, PBS and NBC News—increased.

Most affected sites allowed OpenAI's search crawler, and other providers kept citing many of them, suggesting the shifts reflected ChatGPT's sourcing decisions rather than publisher action.

For ChatGPT, twelve publishers went to (near) zero for a period in 2026




***Note:** some of these publishers, after vanishing from citations, have since seen an increase in citations. Notably APNews and Axios have made it back into ChatGPT's "top cited" list*

Crawler access to a publisher is a stronger predictor for Claude and Gemini than ChatGPT

81% of 66 selected major publishers blocking Claude-SearchBot, received almost no Claude citations, compared with **12%** of publishers that allowed it. For Gemini, **67%** of top publishers blocking Google-Extended received almost no citations, compared with none of the publishers that allowed it.

ChatGPT breaks the pattern: **24%** of publishers allowing OAI-SearchBot still received almost no citations, while some blocked sites continued to appear through other paths.

Share of selected publishers each model rarely cited in Aug 2026, by robots.txt status



AI providers rely on largely different news outlets

The providers share little overlap among their most-cited news sources. Forbes is the only outlet in all four top 10 lists, ranking first for Claude with **19** citations and Gemini with **44**. Perplexity draws from a much deeper news list, led by The New York Times with 47 citations, while ChatGPT's leading outlet, Axios, appears just **12** times. ChatGPT's smaller totals reflect its lower reliance on journalism overall, not concentration among a few outlets.

Which news outlets each provider actually cites

Citations of journalism domains

ChatGPT

axios.com	bbc.co.uk	washingtonpost.com	apnews.com	forbes.com
thepointsguy.com	news.pollstar.com	cbsnews.com	spglobal.com	latimes.com

Claude

forbes.com	cnbc.com	yahoo.com	aol.com	finance.yahoo.com
gulfnews.com	nbcnews.com	foxnews.com	cbsnews.com	consumerreports.org

Gemini

forbes.com	theguardian.com	cnet.com	cbsnews.com	tastingtable.com
simpleflying.com	kiplinger.com	qz.com	businessinsider.com	pcmag.com

Perplexity

nytimes.com	forbes.com	cnbc.com	bbc.com	theguardian.com
finance.yahoo.com	tastingtable.com	businessinsider.com	cbsnews.com	usatoday.com


Claude, Gemini and Perplexity cite much smaller news outlets than ChatGPT

ChatGPT tends to cite newsrooms with big readerships. The median unique visitors monthly, UVM, of ChatGPT citations is over **6.5 million people**.

Meanwhile, the other providers are citing outlets with fewer than **1 million UVMs**.

Claude, in particular, is keen to cite small trade and regional outlets, a pitching opportunity many PR pros miss.

Median monthly unique visitors of the outlets behind journalism citations




Claude more likely to cite pressers from the wire than ChatGPT

Claude and ChatGPT have distinct citation patterns when it comes to press releases. In about **85%** of citations, ChatGPT links to a company's own site—either their newsroom or elsewhere—compared to **58%** of the time with Claude.


Meanwhile, Claude links to press releases on a wire or PR distribution portal in **37%** citations, compared to **14%** with ChatGPT.

Where the AI cited press releases actually live

Hosting of citations classified as press releases—the company's own newsroom vs wires vs republishers



ChatGPT




Claude

Inside Forbes: editorial beats pay-to-play overall

Many PR pros wonder whether Forbes' Councils, their pay-to-play offering, ranks as well as their journalistic content when it comes to GEO with ChatGPT.

While citations practices change depending on the prompt type, in almost all cases, editorial outperforms pay-to-play, although it's not the leading vertical. Affiliate-type content like Forbes Advisor, Vetted, and Home, often outrank them both.

Share of ChatGPT's Forbes.com citations by section,
within each query type




Journalism and owned media serve different roles in GEO

Nearly half of cited journalism, **48%**, links to content about breaking or recent news, while another **30%** of citations point to content like buyer's guides and reviews.


Owned-media citations are dominated by product and service pages at **58%**, followed by practical guides at **24%**.

Journalism is most often cited for timely developments and external evaluation, while brand-owned sites supply product information and guidance.

Content types cited within journalism and owned media



Owned media



Journalism

**What this means
for communicators**

Different questions favor different content formats

No content type performs equally across every kind of question. This table connects six common query intents with the formats AI providers cite most often and the work communicators can do in response. Using these patterns can help teams focus their owned content and earned media efforts on the questions their audiences actually ask.


Content opportunities by query intent

User's question type	Formats cited most often	Opportunity for communicators
Problem discovery	• Practical guides 28% • Product pages 23%	Explain the problem clearly and connect it to a solution
Best-of	• Buyer's guides and reviews 45% • Product pages 26%	Provide verifiable product information and pursue review coverage
Industry trends	• Expert analysis 41% • Recent news 18% • Research 14%	Publish original analysis and offer useful expert commentary
How-to	• Practical guides 68% • Product pages 16%	Create clear, detailed instructional content
Comparative	• Evergreen reference content 37% • Buyer's guides 21%	Develop transparent comparisons and support third-party evaluation
Risk and due diligence	• Practical guides 36% • Evergreen reference content 17% • Expert analysis 14%	Publish evidence-rich guidance and credible expert analysis

No single source strategy reaches every AI provider

No page or placement reliably reaches every provider. This graphic outlines four ways communicators can improve their chances of being cited, from content they control to coverage they earn and distribute.

Using the full mix can broaden visibility across providers and query types while reducing dependence on any single outlet or channel.



Reliable GEO measurement requires scale and consistency

A single response cannot show whether a brand’s visibility is strong or weak. The table translates this study’s findings into a measurement approach that accounts for differences between prompts, providers and repeated runs. Following these practices helps teams separate meaningful changes from the normal variation in AI responses.

Routes to broader AI visibility

Measurement practice	Why it matters
Use a stable prompt set organized by intent	Citation behavior changes depending on the type of question asked.
Run prompts more than once	Providers may cite different sources each time they answer the same question.
Report results by provider	Each provider cites at different rates and relies on a different source mix.
Use ranges and rolling trends	Trends across many responses are more reliable than one-time appearances.
Record model and provider changes	Updates can quickly change citation volume and source selection.
Track citations and answer mentions separately	A brand can appear in an answer without receiving a citation, so the measures capture different forms of visibility.

Thank you!

Generative Pulse helps brands understand and improve how AI portrays them. It's built for brand, marketing, PR, content, SEO and executive teams. Built around the concept of Generative Engine Optimization (GEO) and Answer Engine Optimization (AEO), it shows where brands appear in AI answers, how AI evaluates them and what shapes its recommendations across platforms like ChatGPT and Google AI. Powered by real AI conversations and Muck Rack's media data, it reveals the sources behind those answers and recommends actions across owned content, earned media and website optimization. Generative Pulse is a Muck Rack company. Learn more at generativepulse.ai.

Download the full report and methodology at
<https://generativepulse.ai/report/>