

— THE CITEHAWK AI INDEX

The AI Recommendation Report

What ChatGPT, Claude, Gemini and five more AI models actually recommend when buyers ask, measured across **184** categories and **2,331** ranked brands in the frozen August 2026 edition. They do not agree. That is the story.



The AIs do not agree. Buyers only ask one.

Every day, buyers ask AI models what to buy, which agency to hire, which software to choose. Each AI model answers with conviction, and they answer differently. A brand can be ChatGPT's first recommendation and invisible to Gemini.

This report is built from the CiteHawk AI Index: **1,467** stored answers collected in a single sweep on **August 1, 2026** across 8 AI models, using discovery questions that never name a brand, plus **1,302** regional answers collected August 3 from four countries. Every number traces to a stored answer in the frozen edition.

Three findings

54 of 184

Disagreement: in 54 categories, at least one list-answering AI model leaves the consensus No. 1 out of its answer entirely.

80.8%

Movement: four of five top-10 positions changed hands between the July and August editions.

45 of 46

Regional split: the same question asked from the US, UK, Australia and Canada crowns a different No. 1 in 45 of 46 regionally collected categories.

"AI visibility is not one score. It is eight opinions, revised monthly."

The case for measuring per-model, made throughout this report.

How an AI decides what to recommend

When a buyer asks "best CRM for a small agency", the AI model is not searching. It is recommending: composing an answer from training knowledge, live retrieval, and the sources it trusts. Which brands appear, and in what order, is the new page one.

How we measure it

CiteHawk asks each of the 8 leading AI models the same **discovery questions**: generic buyer questions that never name a brand. An AI model that answers with a brand is recommending it unprompted. Every answer is stored verbatim; a brand enters a ranking only when at least two independent AI models recommend it.

"What are the best running shoes? Recommend the top brands or products that people actually use."

Perplexity · August 2026 edition · stored verbatim

The best running shoes depend on your use case, but the brands and models that keep showing up in 2026 expert picks are **Adidas, Brooks, New Balance, Saucony, ASICS, Nike, and Hoka**. [1][5] [12][14] If you want the most practical "people actually use them" shortlist, these are the strongest mainstream options: Best overall / popular all-around pick: **Adidas Adizero Evo SL** [1][5] ...

The anatomy: a discovery question in, a conviction-filled answer out. Seven brands surfaced, in the AI model's own order, with the sources it cited. This answer is one of the 1,467 underpinning this edition.

ChatGPT

184 answers

Claude

184 answers

Gemini

179 answers

Perplexity

184 answers

DeepSeek

184 answers

Grok

184 answers

Copilot

184 answers

Google AI

184 answers

Answer counts for the August 2026 edition. Five Gemini responses failed collection and are absent from the frozen record; Google AI Overviews returned a brand-named answer in 33 of its 184 responses (see Section 04).

Eight AI models, eight opinions

Winners are real: in every one of the 184 categories, the No. 1 is named by at least four AI models. The disagreement lives everywhere below that line.

RUNNING SHOES · TOP 8	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	DEEPSEEK	GROK	COPILOT	GOOGLE AI
Asics								
Saucony								
Adidas								
Nike								
Hoka								
New Balance								
Brooks								
Puma								

Filled cell = the AI model named the brand in its August 2026 answer. Google AI returned no brand-named overview for this question; the blank column is the finding.

The big four agree, and then the field splits: Claude and Gemini skip **Brooks**, five AI models skip **Puma**, and Copilot leaves out **Hoka** and **New Balance**. In 54 of the 184 categories, at least one list-answering AI model omits the consensus No. 1 itself; in only 130 do all seven name it.

"Below the winner, agreement is the exception."

MOST CONTESTED CATEGORIES	SCORE	MOST CONSENSUS	SCORE
Online MBA programs	2.7	Online course platforms	6.2
Wireless earbuds	2.7	Travel apps	6.1
Smart thermostats	2.8	HVAC companies	6.0
Smart locks	3.0	Job search sites	5.9
Business credit cards	3.1	Accounting software	5.9

Score = mean AI models (of eight) naming each of the category's top-10 brands, August 2026 edition. Low = contested, high = consensus.

The rankings moved while you were reading this

August is the Index's first month-over-month edition, and the first time movement is visible at scale: **80.8%** of top-10 positions changed hands in one month.

CATEGORY	BRAND	JULY	AUGUST	CHANGE
Credit cards	Chase Freedom Unlimited	15	1	▲ 14
Water filters	Waterdrop	14	1	▲ 13
Wireless earbuds	Sony WF-1000XM6	12	1	▲ 11
Pillows	Purple	13	2	▲ 11
Electric SUVs	Rivian	12	2	▲ 10
Graphic design software	Adobe Creative Cloud	13	3	▲ 10
Travel apps	KAYAK	3	out	▼
Coffee subscriptions	Blue Bottle Coffee	3	out	▼
Credit cards	Capital One Venture Rewards	4	out	▼
Dog food	Hill's Science Diet	4	out	▼

Positions in each category's full consensus ranking, July and August 2026 editions. "Out" = the brand no longer meets the two-model consensus gate in August. Churn: of 1,730 comparable top-10 slots, 1,397 changed occupant.


One brand's month: Sony WF-1000XM6, wireless earbuds

Claude 1 → 1 held No. 1	Copilot — → 1 absent in July	Gemini — → 1 absent in July	Grok — → 1 absent in July
--	---	--	--

In July, Claude alone ranked it first. By August, four AI models call it the best wireless earbuds, and that is enough to make it the category's consensus No. 1, named by four of eight. AI models retrain, retrieve fresher sources, and re-weigh what they trust. A ranking is a snapshot, never an asset you own.

Every AI model has a personality

Per-model splits reveal systematic habits: some hedge with long lists, some concentrate on the consensus winners, and one usually declines to name brands at all.



Left: mean brands surfaced per discovery answer, August 2026 edition. Right: of all brand mentions each AI model made, the share going to the category's consensus No. 1. Both computed from the 1,467 stored answers.

Grok hedges hardest, spreading sixteen names across an answer and giving the consensus winner the smallest share. Perplexity concentrates: fewest names among the list-answering seven, largest share to the winner. And Google AI Overviews mostly abstains: it returned a brand-named answer in 33 of 184 categories, and its pick matched the consensus No. 1 in just one of those.

What this means when you read your own numbers

A gap on a hedging AI model is easier to close than a gap on a committed one; sixteen names means room on the list. Per-model strategy beats a single blended score, and that is why CiteHawk reports all 8 separately.

"Ask Grok for the best in a category and you get sixteen names. Ask Google AI and, four times out of five, you get none."

Ask from Sydney, get a different answer

The August edition adds regional collection: the same discovery questions asked from the US, UK, Australia and Canada.

"What are the best business bank accounts?" · asked from four countries, August 3

United States

Mercury

then Bluevine · Chase
Business Complete
Banking

United Kingdom

Tide

then Barclays · Monzo
Business

Australia

NAB

then Westpac ·
Commonwealth Bank

Canada

TD

then Scotiabank ·
CIBC

The consensus No. 1 per country in the August 3 regional sweep, runners-up beneath. The global edition crowns Relay. Four countries and the global edition produce five different No. 1s.

45 of 46

regionally collected
categories crown a different
No. 1 in at least two of the
four countries.

37

of those categories crown
four different No. 1s, one
per country.

1,302

regional answers collected
August 3 and frozen into
the edition.

For global brands this is a blind spot inside a blind spot: visibility earned in one market silently fails to transfer. Of the 46 regionally collected categories that produce a consensus ranking in every country, 45 crown a different No. 1 in at least two of them. In 37, all four countries disagree, and in 45 at least one country's answer differs from the global edition's.

Regional collection covers the 47 region-flagged categories of the Index, the ones where local buyers get local answers, asked with the buyer placed in each country; 46 of the 47 produced a consensus ranking in all four. Same discovery questions, same 2-model consensus gate, one run per AI model. Gemini returned no usable answers in the regional sweep and contributes nothing to these numbers.

What a brand does about it

The pattern across every finding: AI visibility is per-model, perishable, and regional. The playbook follows directly.

1 · Measure per-model, on a rhythm

A blended score hides the gap that is costing you buyers. Weekly workspace tracking matches how fast the answers actually change; the Index's monthly editions freeze the record so movement is provable.

2 · Fix what the AIs get wrong first

AI models repeat wrong facts confidently. Correcting the record at the sources they cite is the fastest, cheapest visibility win.

3 · Win the sources, not the algorithm

Every recommendation traces to sources the AI model trusts. Presence on the directories, reviews and publications an AI model cites is the durable lever.

4 · Treat rankings as a feed, not a trophy

When 80.8% of top-10 positions change hands in a month, the work is a rhythm: review weekly, act monthly, verify continuously.

Methodology & integrity

Collection

Discovery questions (never naming a brand) asked monthly across 8 AI models at consumer parity: the same models consumers use, not APIs tuned differently. All answers stored verbatim; 1,467 global and 1,302 regional answers underpin this edition.

Ranking

Presence-based; a brand enters a ranking only when recommended by at least 2 independent AI models. Full methodology, including scoring and tie rules: citehawk.com/leaderboards/methodology.

Integrity

Rankings cannot be bought. Claims are human-reviewed and never affect position. Badges are dated awards frozen to the edition that earned them; a living certificate adds each edition a brand places in and never claims a live rank. Frozen monthly editions are permanent records.

Edition facts

184

published categories
(August 2026 edition)

2,331

brands ranked across all
categories

1,467

stored answers behind this
edition

8

AI models at consumer
parity

Aug 1

global sweep, 04:04 to
05:01 UTC

4

regions collected August 3:
US, UK, AU, CA (1,302
answers)

Five Gemini responses failed the global collection, and Gemini returned no usable answers in the regional sweep; failed answers are absent from the frozen record, and every statistic in this report is computed from the answers that exist, never imputed.

— YOUR BRAND, IN THESE ANSWERS

Find out what the AIs say when buyers ask about your category.

CiteHawk runs your brand through the same 8 AI models with a free 50-prompt visibility report. Every answer stored, every gap named, ready within the hour.

Run your free report → citehawk.com

© 2026 ScaleUp Teams Pty Ltd · The CiteHawk AI Index · Data from the frozen August 2026 and July 2026 editions · Reproduction with attribution welcome