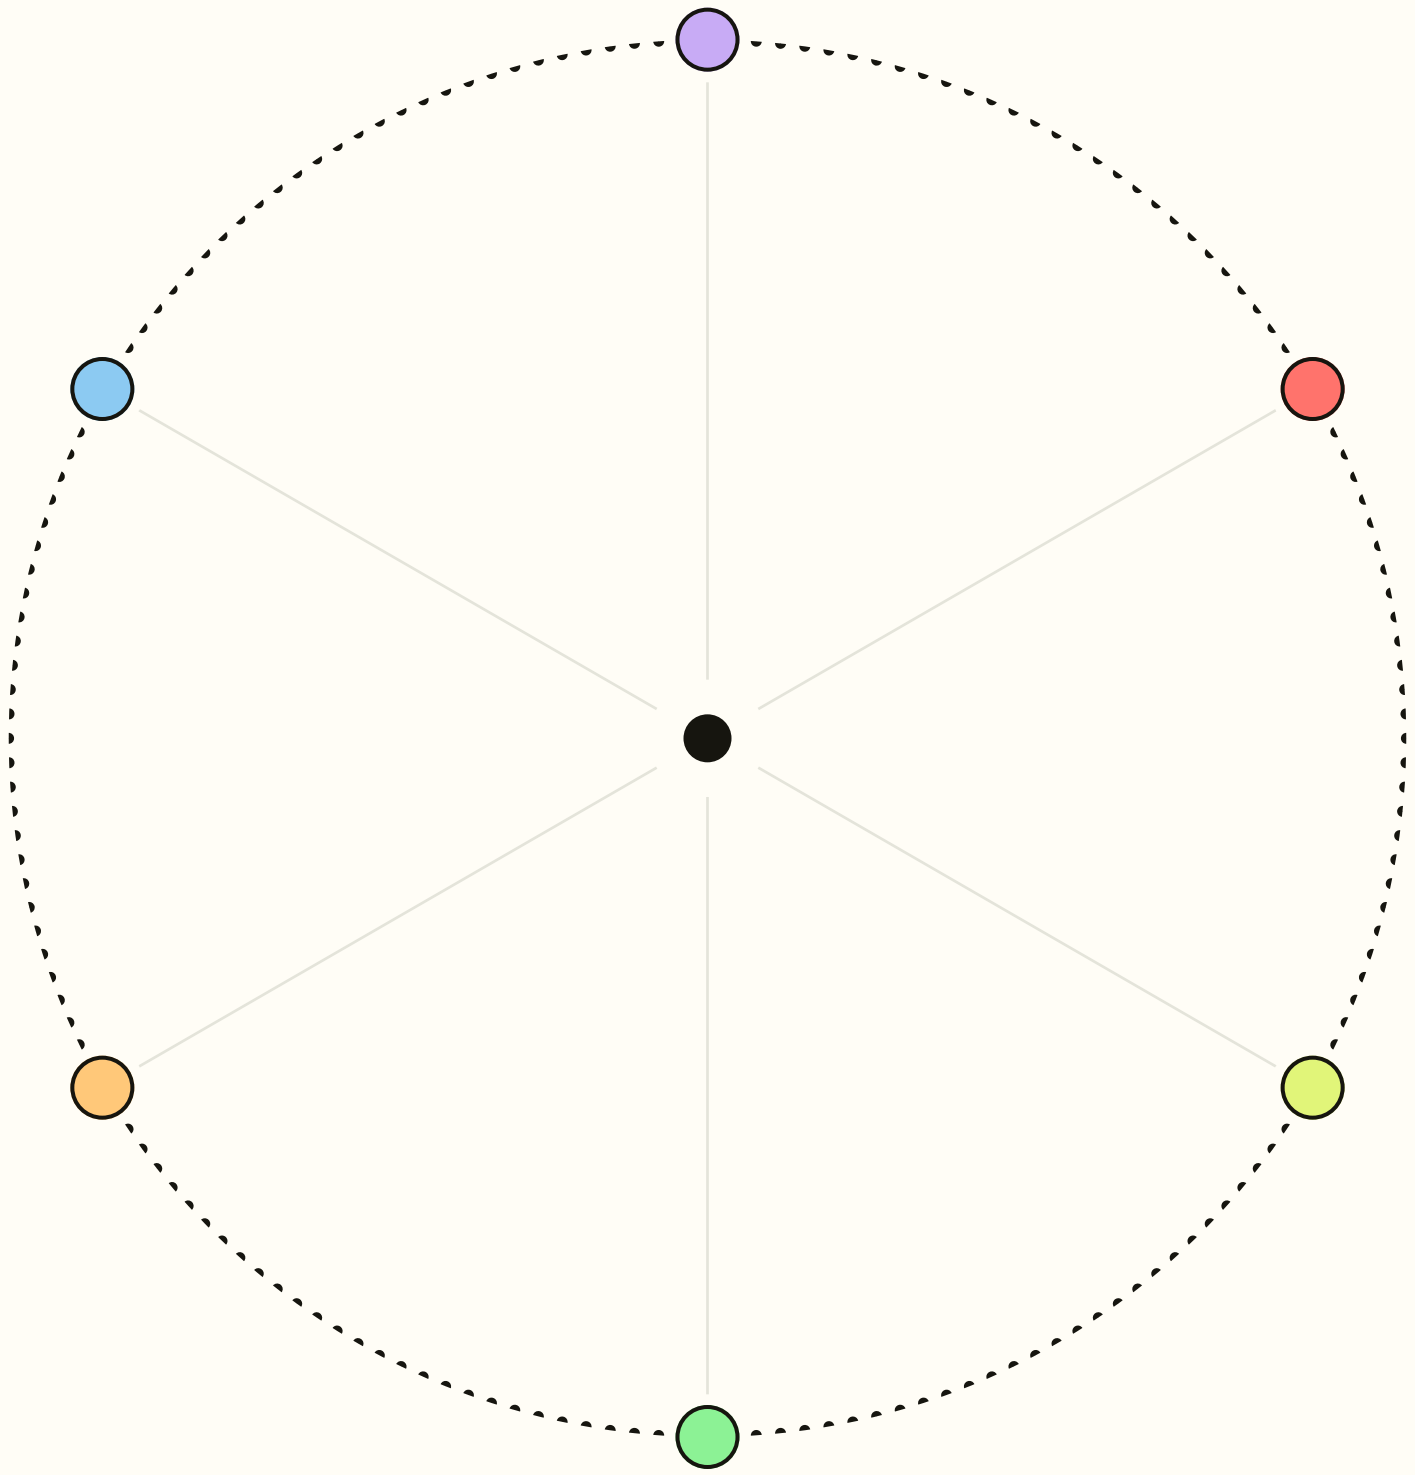


The AI Search Operating Manual.

How to be found, trusted, cited, and chosen across AI search.

ATLAS

by Momentic



The AI Search Operating Manual.

How to be found, trusted, cited, and chosen across AI search & discovery surfaces.

Contents.

01 Why this exists 03

02 How AI search works 06

03 The four gates 12

04 The three forces 14

05 The centroid 18

06 The Influenceable Web 22

07 Demand 26

08 Content that survives 30

09 Technical eligibility 34

10 Platform punchlists 38

11 Commerce, local, agents 42

12 Measurement 46

13 The 90-day track 50

14 The audit 53

15 The myths 55

16 The one-page version 58



01

Why this exists.

What Atlas measurement has found to date.

Atlas ran 8,209 prompts across 11 AI engines for 16 brands from July to September 2026. These six findings recur through the manual.

17.3% The same prompt on the same engine changed whether the brand was mentioned across repeated runs. p.11

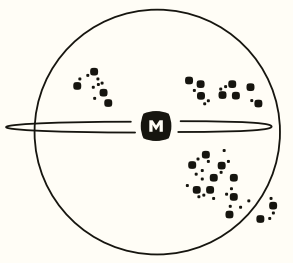
2 to 5 AI engines that show their searches ran a median of 2 to 5 per answer. p.10

89.3% of citations pointed to websites the brand does not own. p.24

5.5x Brands were mentioned 5.5 times as often as they were recommended on category prompts. p.48

20x Owned articles updated in the last 90 days were cited 20 times as often per page as articles older than a year. p.15

3x Firewalls and CDNs refused AI crawlers three times as often as robots.txt blocked them. p.36



Source: Atlas measurement, 16 brands, 11 engines, 2026-07-13 to 2026-09-14. Each page states its finding's exact window. Brands not named.

How much to trust each recommendation.

Every recommendation in this manual is tagged with one of these four tiers. The tier decides what the recommendation is allowed to claim.

TIER	WHAT IT MEANS	HOW TO TREAT IT
01 **Platform-confirmed**	The provider documents the behavior or the control.	Act on it, and re-check the date because products change.
02 **Research**	A controlled study changed a variable and measured the outcome.	Apply it only to the conditions the study tested.
03 **Observational**	A large dataset shows an association.	Use it to decide what to test first.
04 **Hypothesis**	Plausible, repeatedly observed, unproven.	Test it before recommending it to a client.



Momentic studio portrait.

Tactics with no evidence at any tier are listed in the myth section instead of recommended.

PART 15, THE MYTHS, PAGE 55

Evidence tiers, defined in part 01 of this manual.

02

How AI search works.

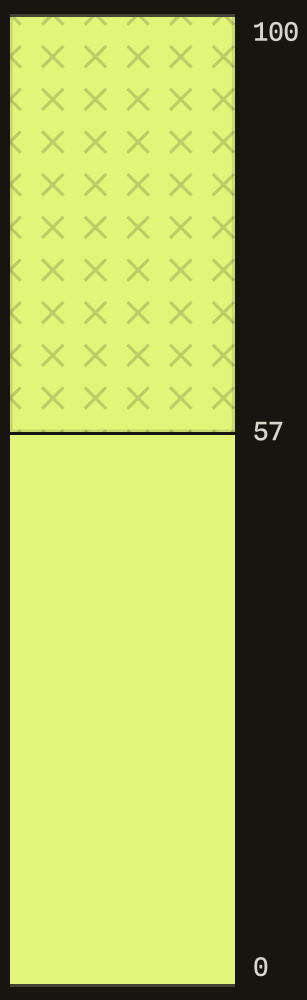
Brands missing from training data start 57 points behind.

A brand absent from training data starts 57 points behind, and no on-page work closes that gap.

57 / 100

RERANKER SCALE, 0 TO 100

Published tests of invented brands against real ones report roughly a 57-point gap.



57

points, the reported difference between an invented brand and a real brand before any on-page work.

Only the difference is reported, so the segment is anchored at zero. The absolute scores are not published.

Source: published tests of invented brands against real brands on a 100-point reranker scale.

Search and AI answers share the same first 6 steps.

The same six stages run before either product renders anything, and a page can be excluded at any one of them.

01

Query understanding

The system converts the prompt to an embedding before it retrieves anything.

02

Query transformation

The system rewrites the prompt and expands it into many search variants.

03

Multi-path retrieval

The system retrieves candidates along more than one path at once.

04

Candidate processing

The system converts retrieved items into candidates it can compare.

05

Multi-stage ranking

Cheap rankers run first, and expensive rankers run only on the survivors.

06

Eligibility and filtering

Filters remove candidates the product is not allowed to use.

The search path and the grounding path diverge after stage 06.

After that, search and AI answers split.

After filtering, the same candidate pool feeds both products, and each product applies its own selection rules.

ONE FILTERED CANDIDATE POOL

Query understanding through eligibility and filtering

THE BRANCH POINT

SEARCH PATH

Composes a results page

Links

Modules

Snippets

GROUNDING PATH

Constructs an answer with citations

Selects evidence

Constructs passages with attribution

Generates a constrained answer

Cross-checks the answer

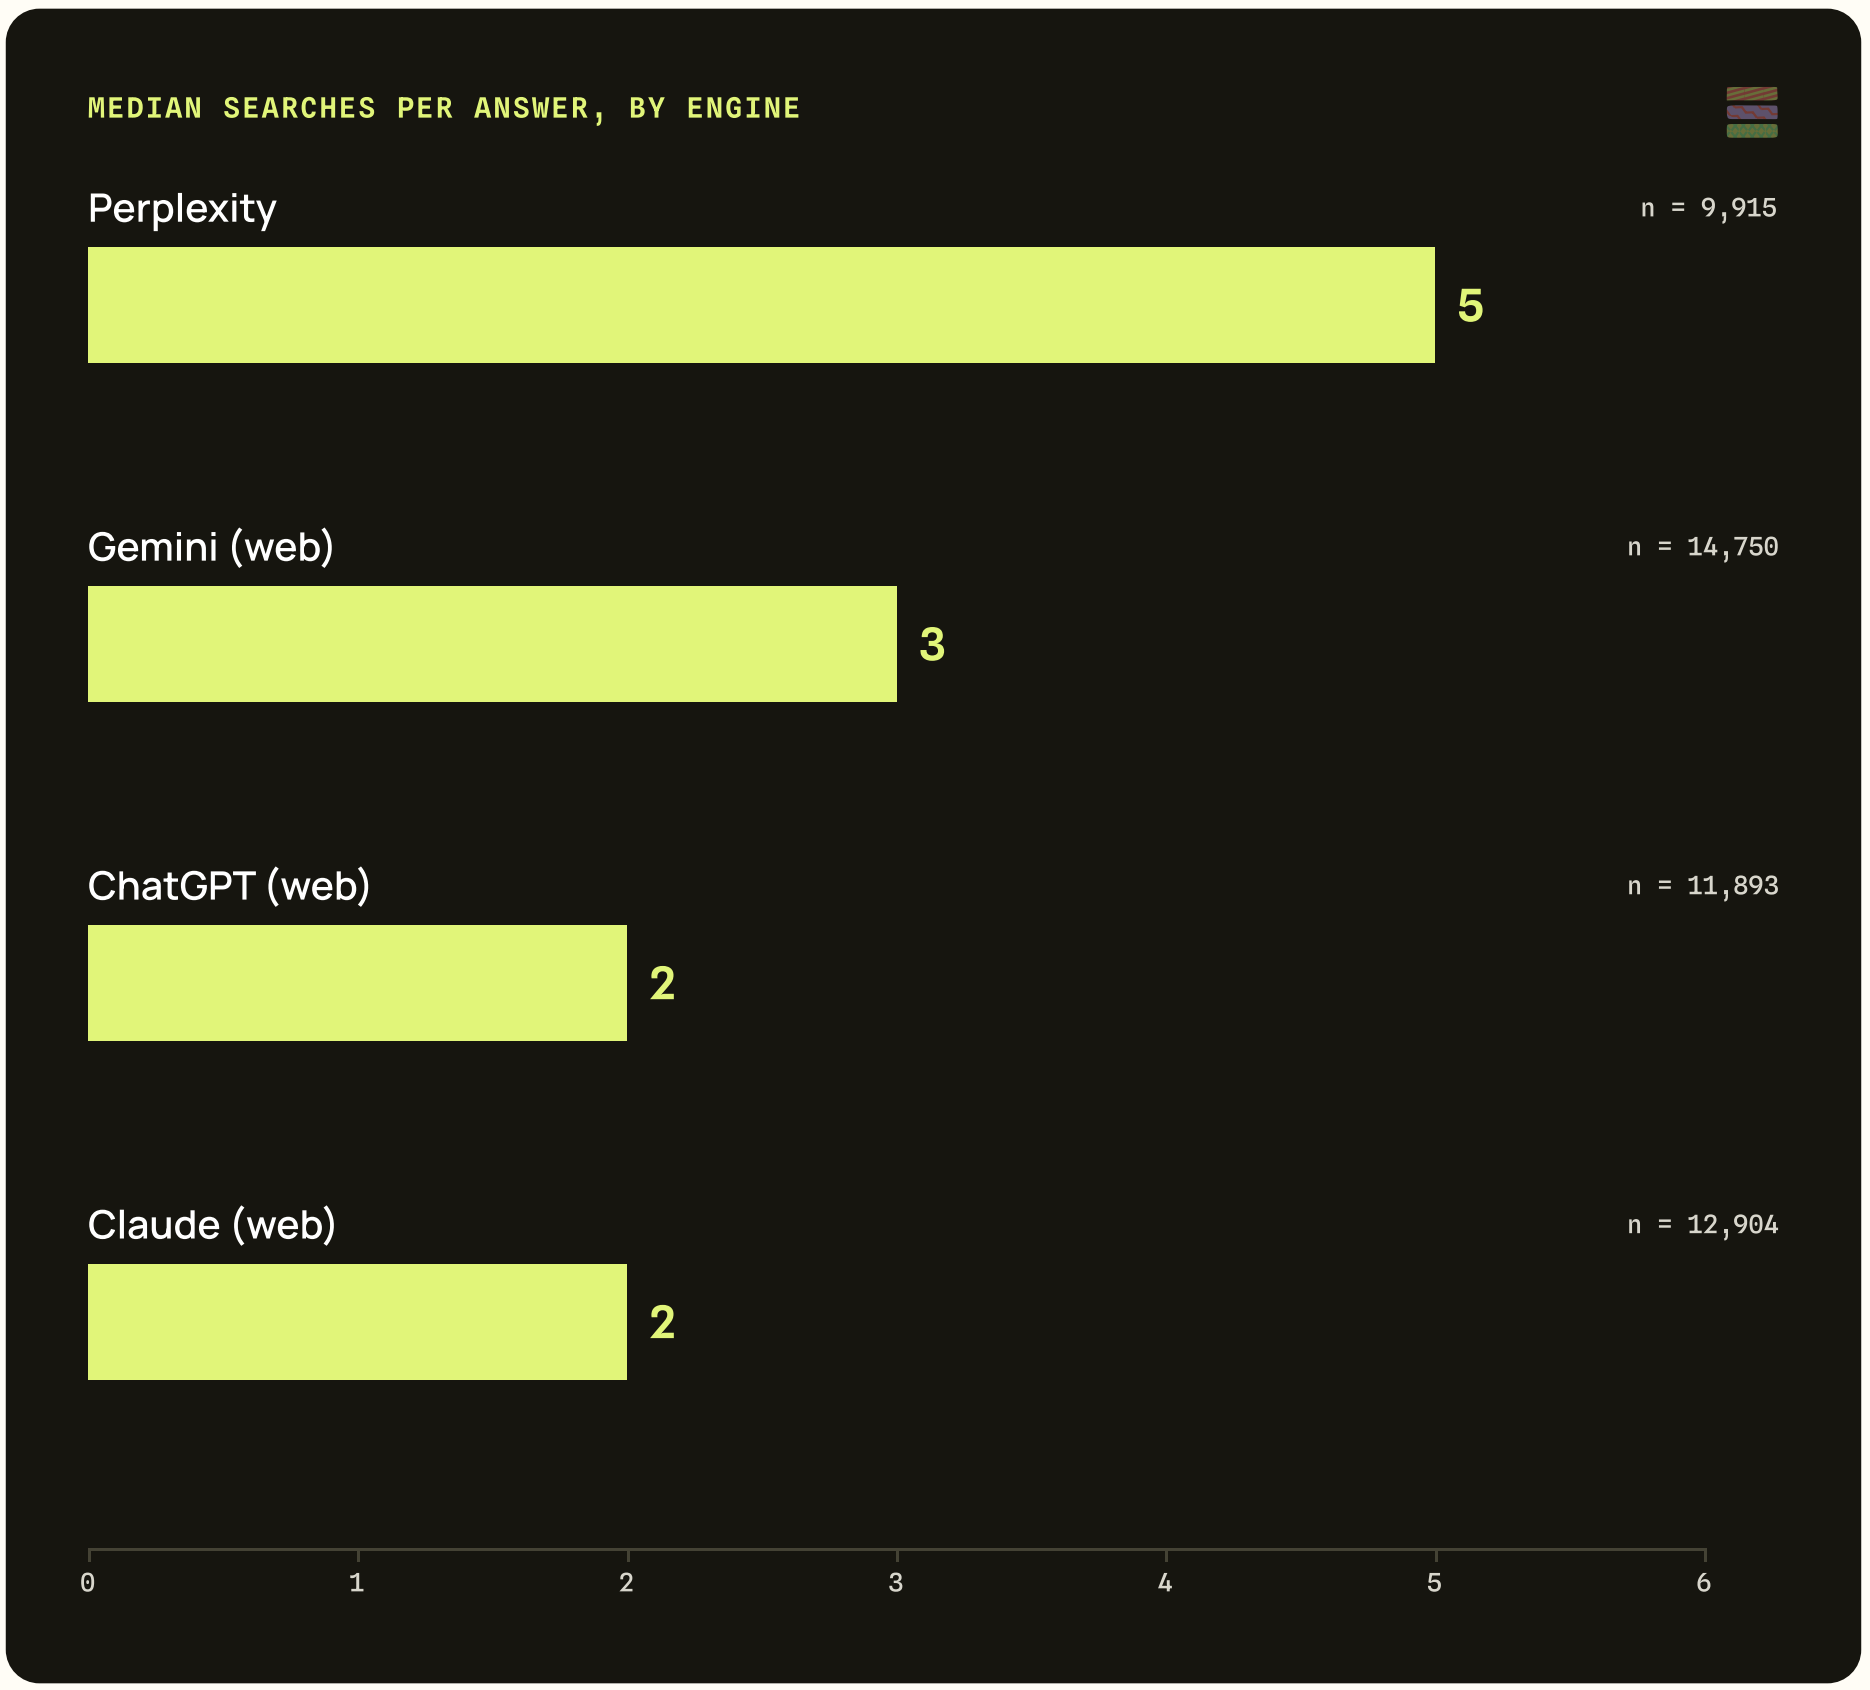
Attaches citations

A page tuned for the results page can be excluded from evidence selection, and the reverse.

AI engines run 2 to 5 searches per answer.

When an engine shows the searches it ran, Atlas records them. Perplexity ran a median of 5 per answer; ChatGPT and Claude with web search ran 2.

This manual previously cited about 15 search variants and 600 candidates per prompt. Atlas could not verify that figure, so the bars below show only the searches engines exposed.

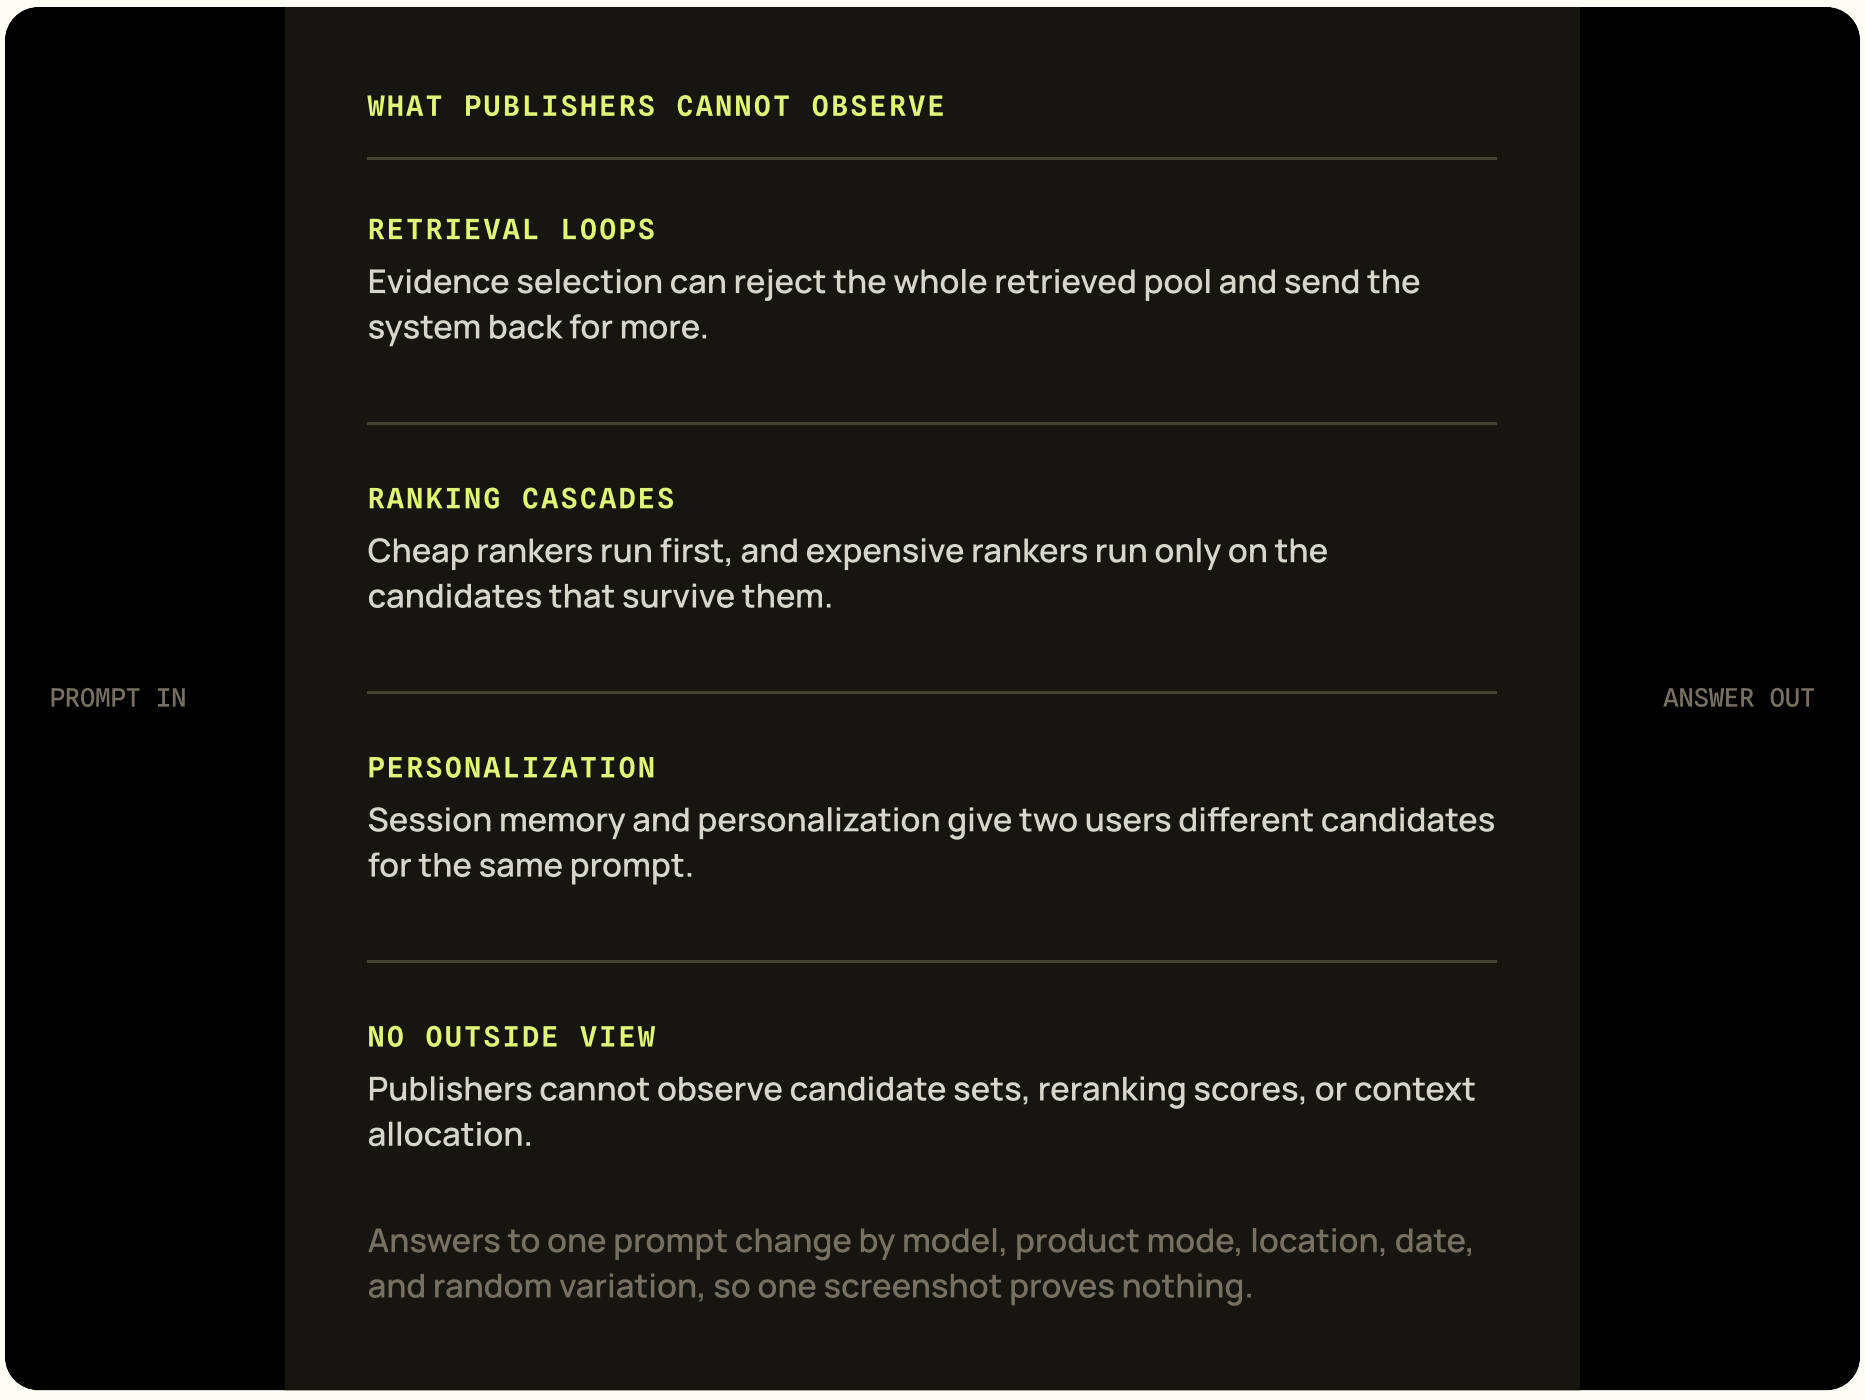


A page written for one phrasing matches only one of those searches.

Source: Atlas measurement, 16 brands, 4 engines that expose their searches, 49,462 answers, 2026-07-27 to 2026-09-14. The other 7 engines do not expose searches.

You can’t fully observe why and how AI systems rank your pages.

Publishers cannot measure the ranking stages directly, so find the gate a page does not pass before applying any checklist.



ANSWERS VARY BETWEEN RUNS

17.3% of prompt and engine pairs changed whether the brand was mentioned across repeated runs.

Google AI Mode changed 0.1% of the time within one run; the Claude and Gemini APIs changed 10%.

Source: Atlas measurement, 16 brands, 11 engines, 29,563 prompt and engine pairs with 2 or more runs, 2026-07-13 to 2026-09-14.

03

The four gates.

A page must pass four checks to be cited.

A page that does not pass one gate is excluded from the answer, regardless of the other three.

01 / 04

01

Discoverable

Can the system reach it?

Crawlable URL, sitemap, no robots block, clean canonical. Also check for a firewall or CDN that challenges the bot by IP while robots.txt allows it.

02

Parsable

Can the system extract structure?

Semantic HTML, valid structured data, alt text, transcripts, clean heading hierarchy. A visually impressive JavaScript blob with no semantic structure does not pass this gate.

03

Groundable

Can a passage be trusted on its own?

Self-contained claims, named sources, no contradictions on the page. Grounding stages deprioritize ambiguous passages.

04

Trustworthy

Do other sources and time agree?

Entity data that matches public knowledge graphs, attributed quotes, and the same claim stated on several credible websites over time.

A gate is pass or fail, never a score.

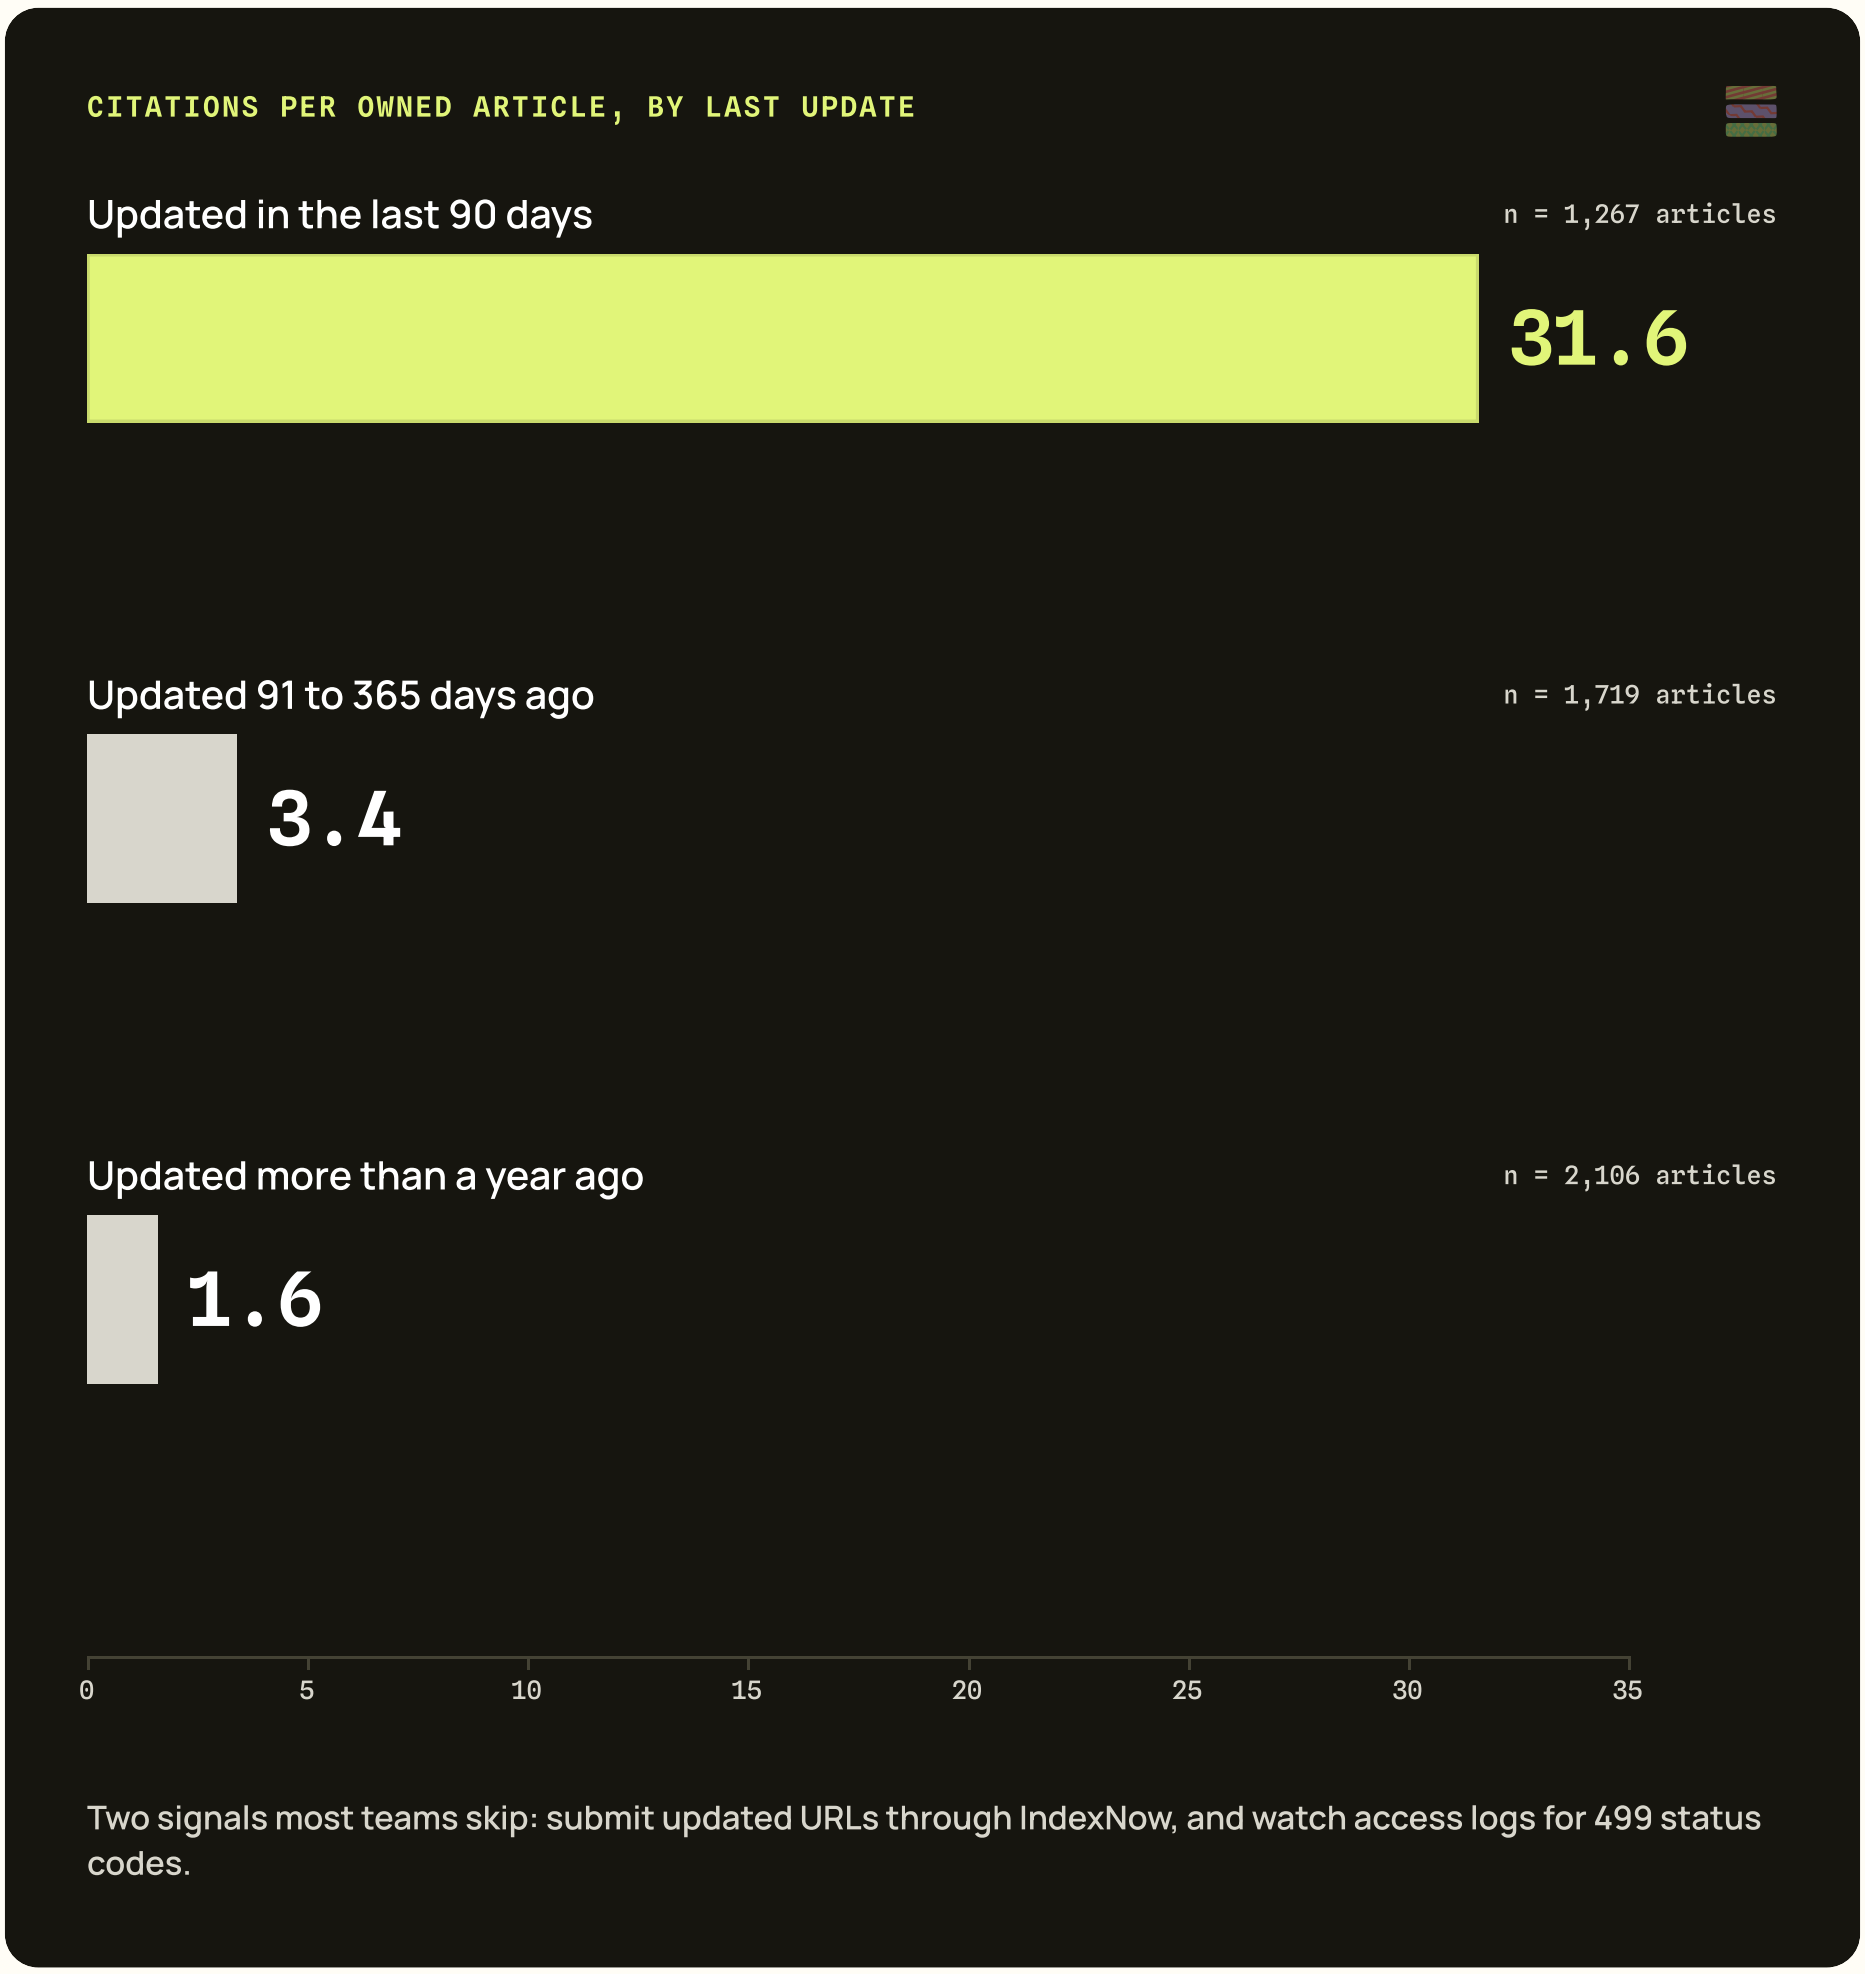
Part 03, the four gates

04

The three forces.

Recently updated articles were cited far more often.

Atlas matched owned articles to the citations they received. Articles updated in the last 90 days were cited about 20 times as often per page as articles older than a year.



Source: Atlas measurement, 13 brands, 5,092 owned articles with a published modified date, 2026-07-13 to 2026-09-14. Observational: frequently updated pages may also be the pages brands consider most important.

Cite your sources so AI systems can verify your claims.

Verification stages check named sources, specific data, and traceable claims against each other. A stage can drop a whole passage when one of its claims is contradicted.

> **“Studies show” is unverifiable. A named source with a year is evidence.**

AI answers quote single passages from indexed pages.

Retrieval selects whole pages. Answer generation then quotes one passage from a selected page.

4 SELECTED / 4 SKIPPED

PASSAGES THAT GET SELECTED	PASSAGES THAT GET SKIPPED
01 **A direct answer in the first sentences**	01 **Background prose**
02 **Definition blocks**	02 **Methodology preambles**
03 **Evidence-dense data**	03 **Transitions**
04 **Numbered steps**	04 **Marketing copy**

Models weight the start and end of a passage most, so an introduction placed before the answer lowers its chance of selection. Put the answer first and the evidence after it.



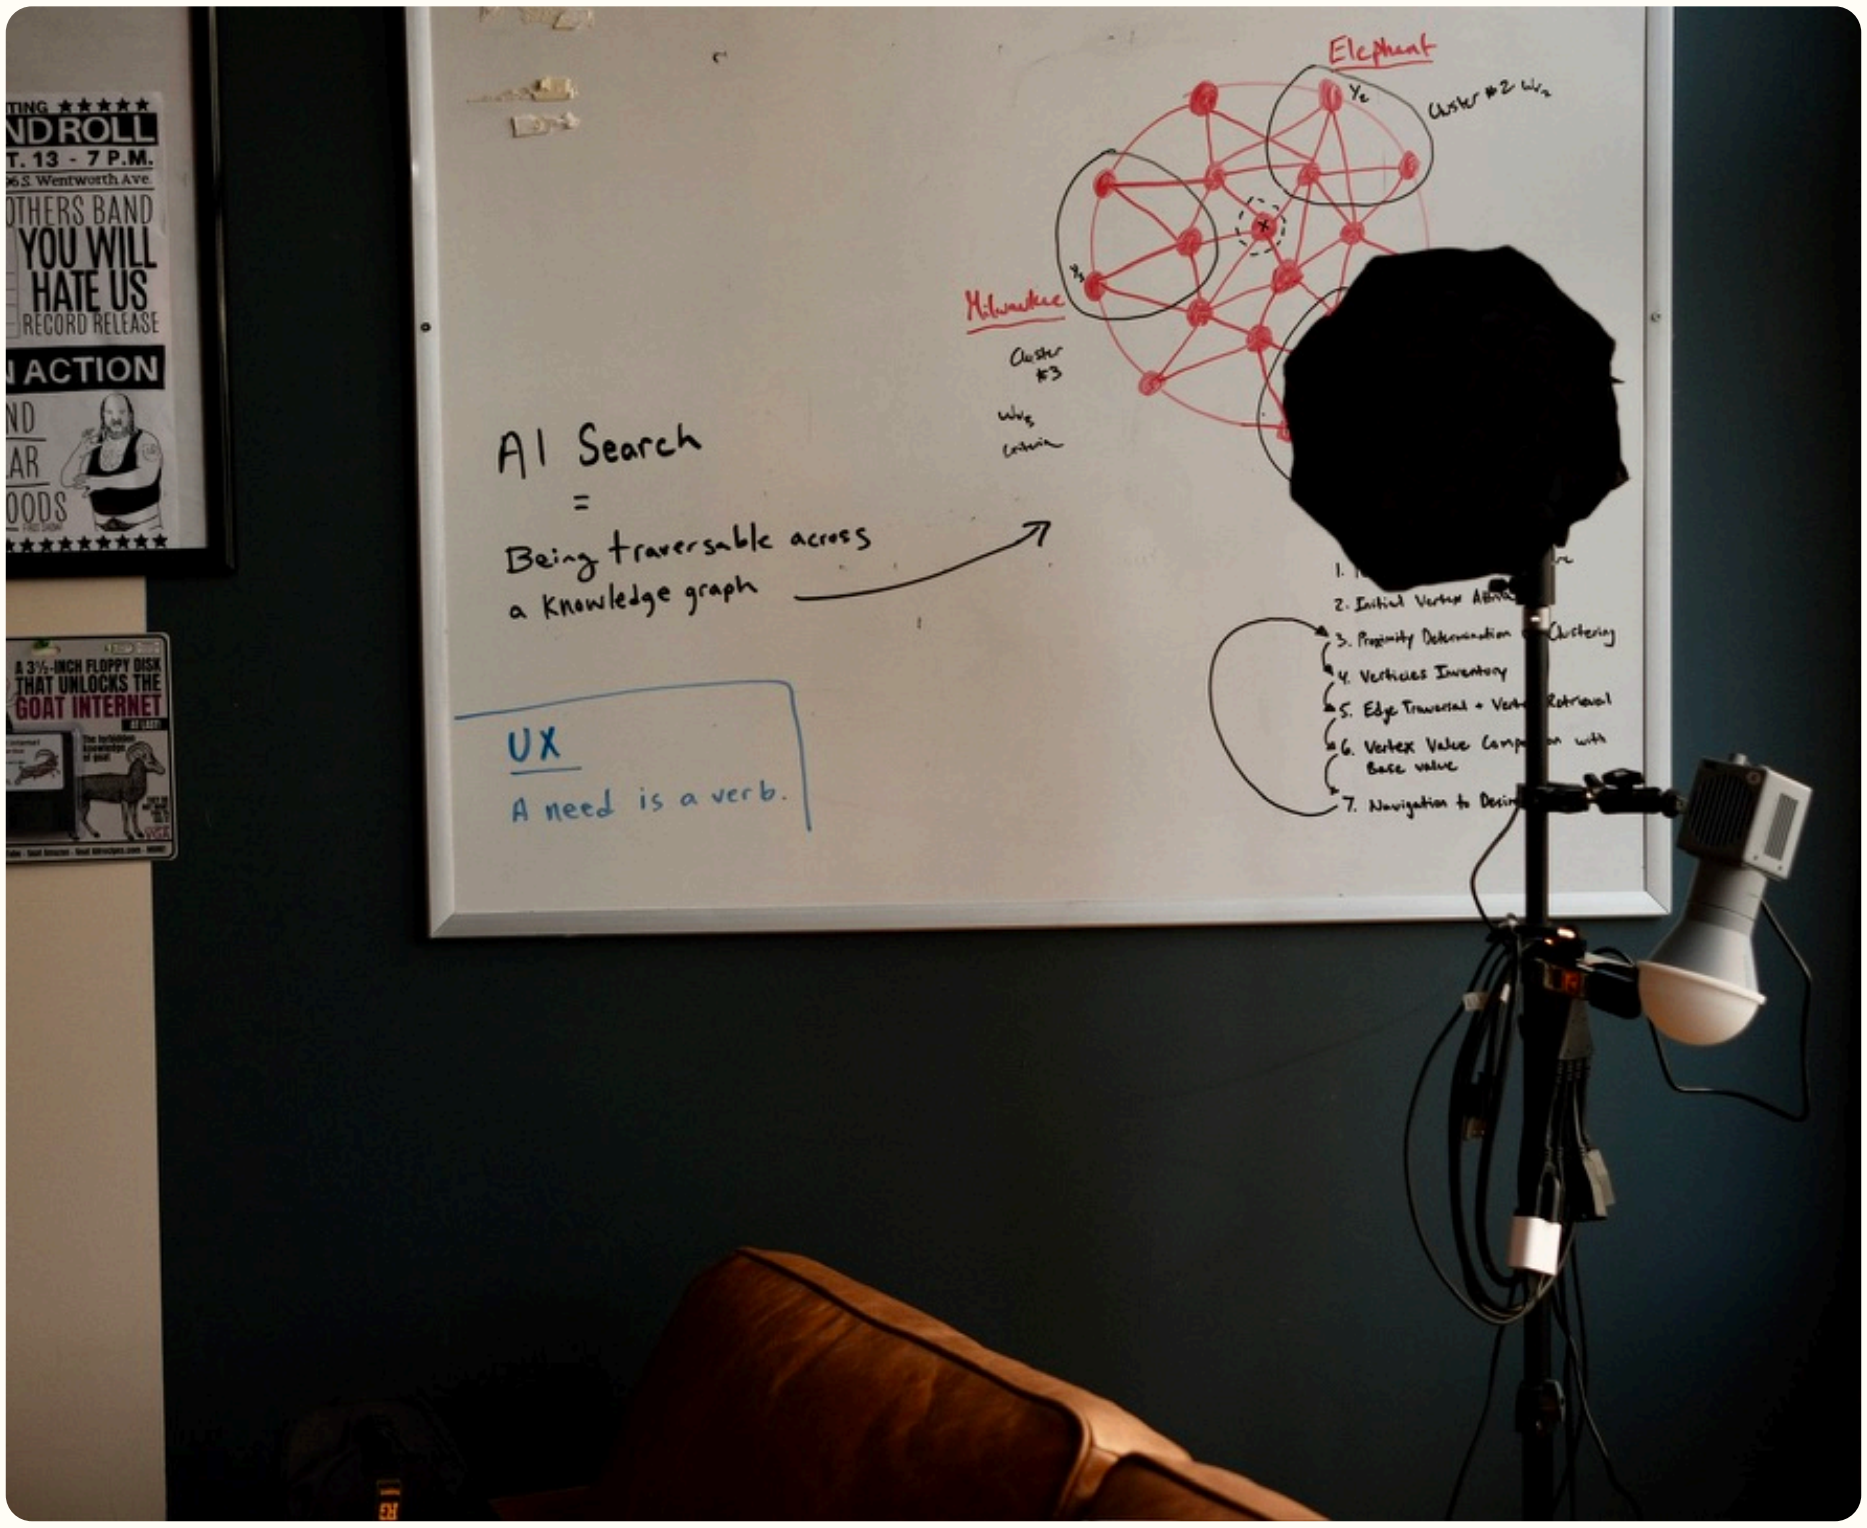
A direct answer is good for people and for AI.

05

The centroid.

A brand is the average of everything it publishes, and everything public said about it.

In vector space that average is called a centroid. Models build their picture of your brand from that text.



Momentic studio wall.

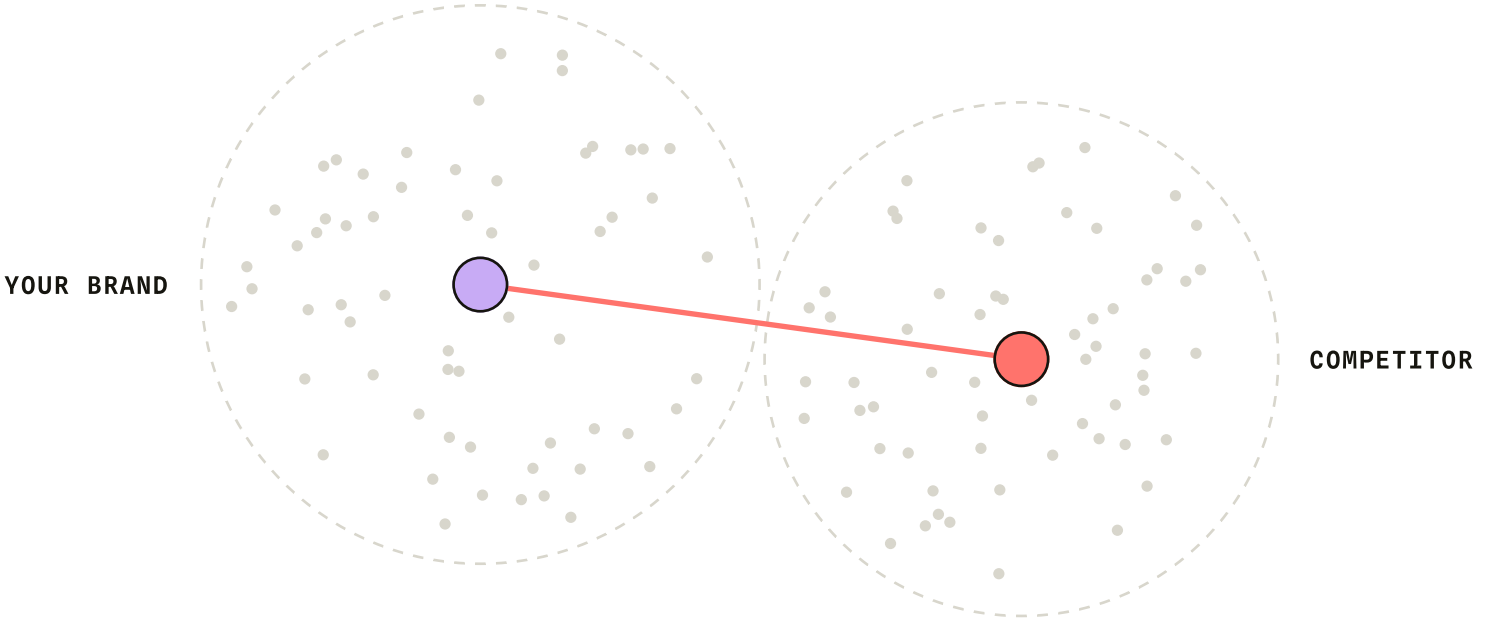
Each page is one point. A brand cited across many prompts covers a connected region of related topics.

How brands become interchangeable.

A centroid far from competitors' centroids makes the brand easy to tell apart. Two failure modes erode that: cluster collision and centroid drift.

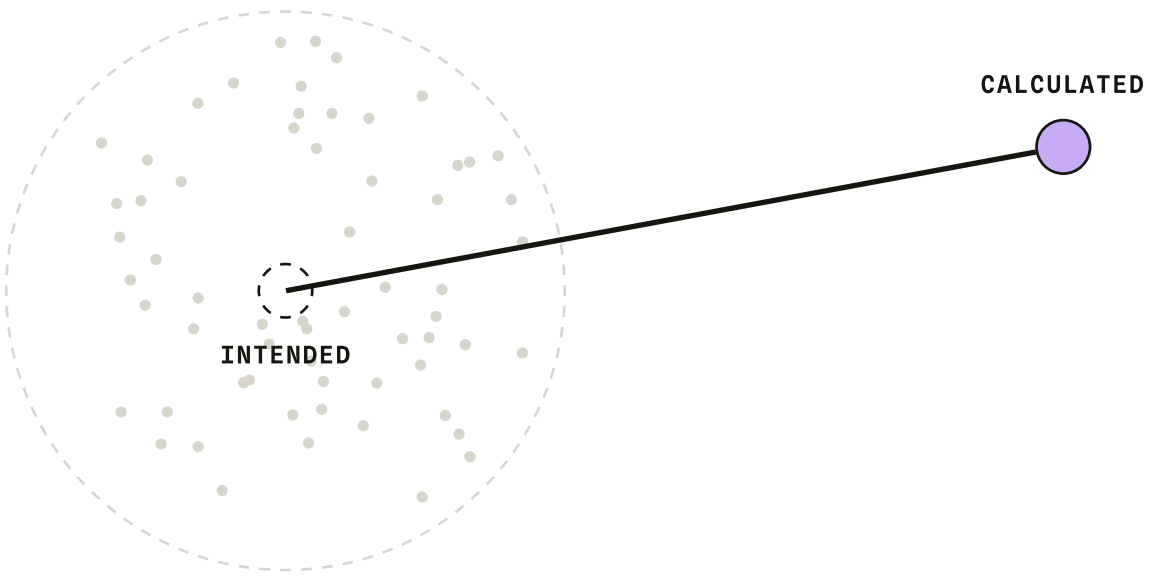
CLUSTER COLLISION

When a brand uses the same sources, templates, and talking points as competitors, models treat the brands as interchangeable.



CENTROID DRIFT

Off-brand campaigns and filler content shift the calculated brand away from how the brand describes itself.



Each small dot is one piece of published content. The large dot is the average.

THIS IS SCHEMATIC AND NOT PLOTTED FROM MEASURED DATA.

How AI selection criteria have changed.

The criterion that decides selection changed twice: presence first, then similarity to the query, now connectedness of evidence.

ERA	WHAT MATTERED	THE WORK
01 **Tokens**	Presence in training corpora	Get the brand into the sources models train on.
02 **Embeddings**	Similarity to the query	Publish pages that match the questions buyers ask.
03 **Traversal** CURRENT ERA	Connectedness of evidence	Connect entities, relationships, and third-party confirmation across websites.



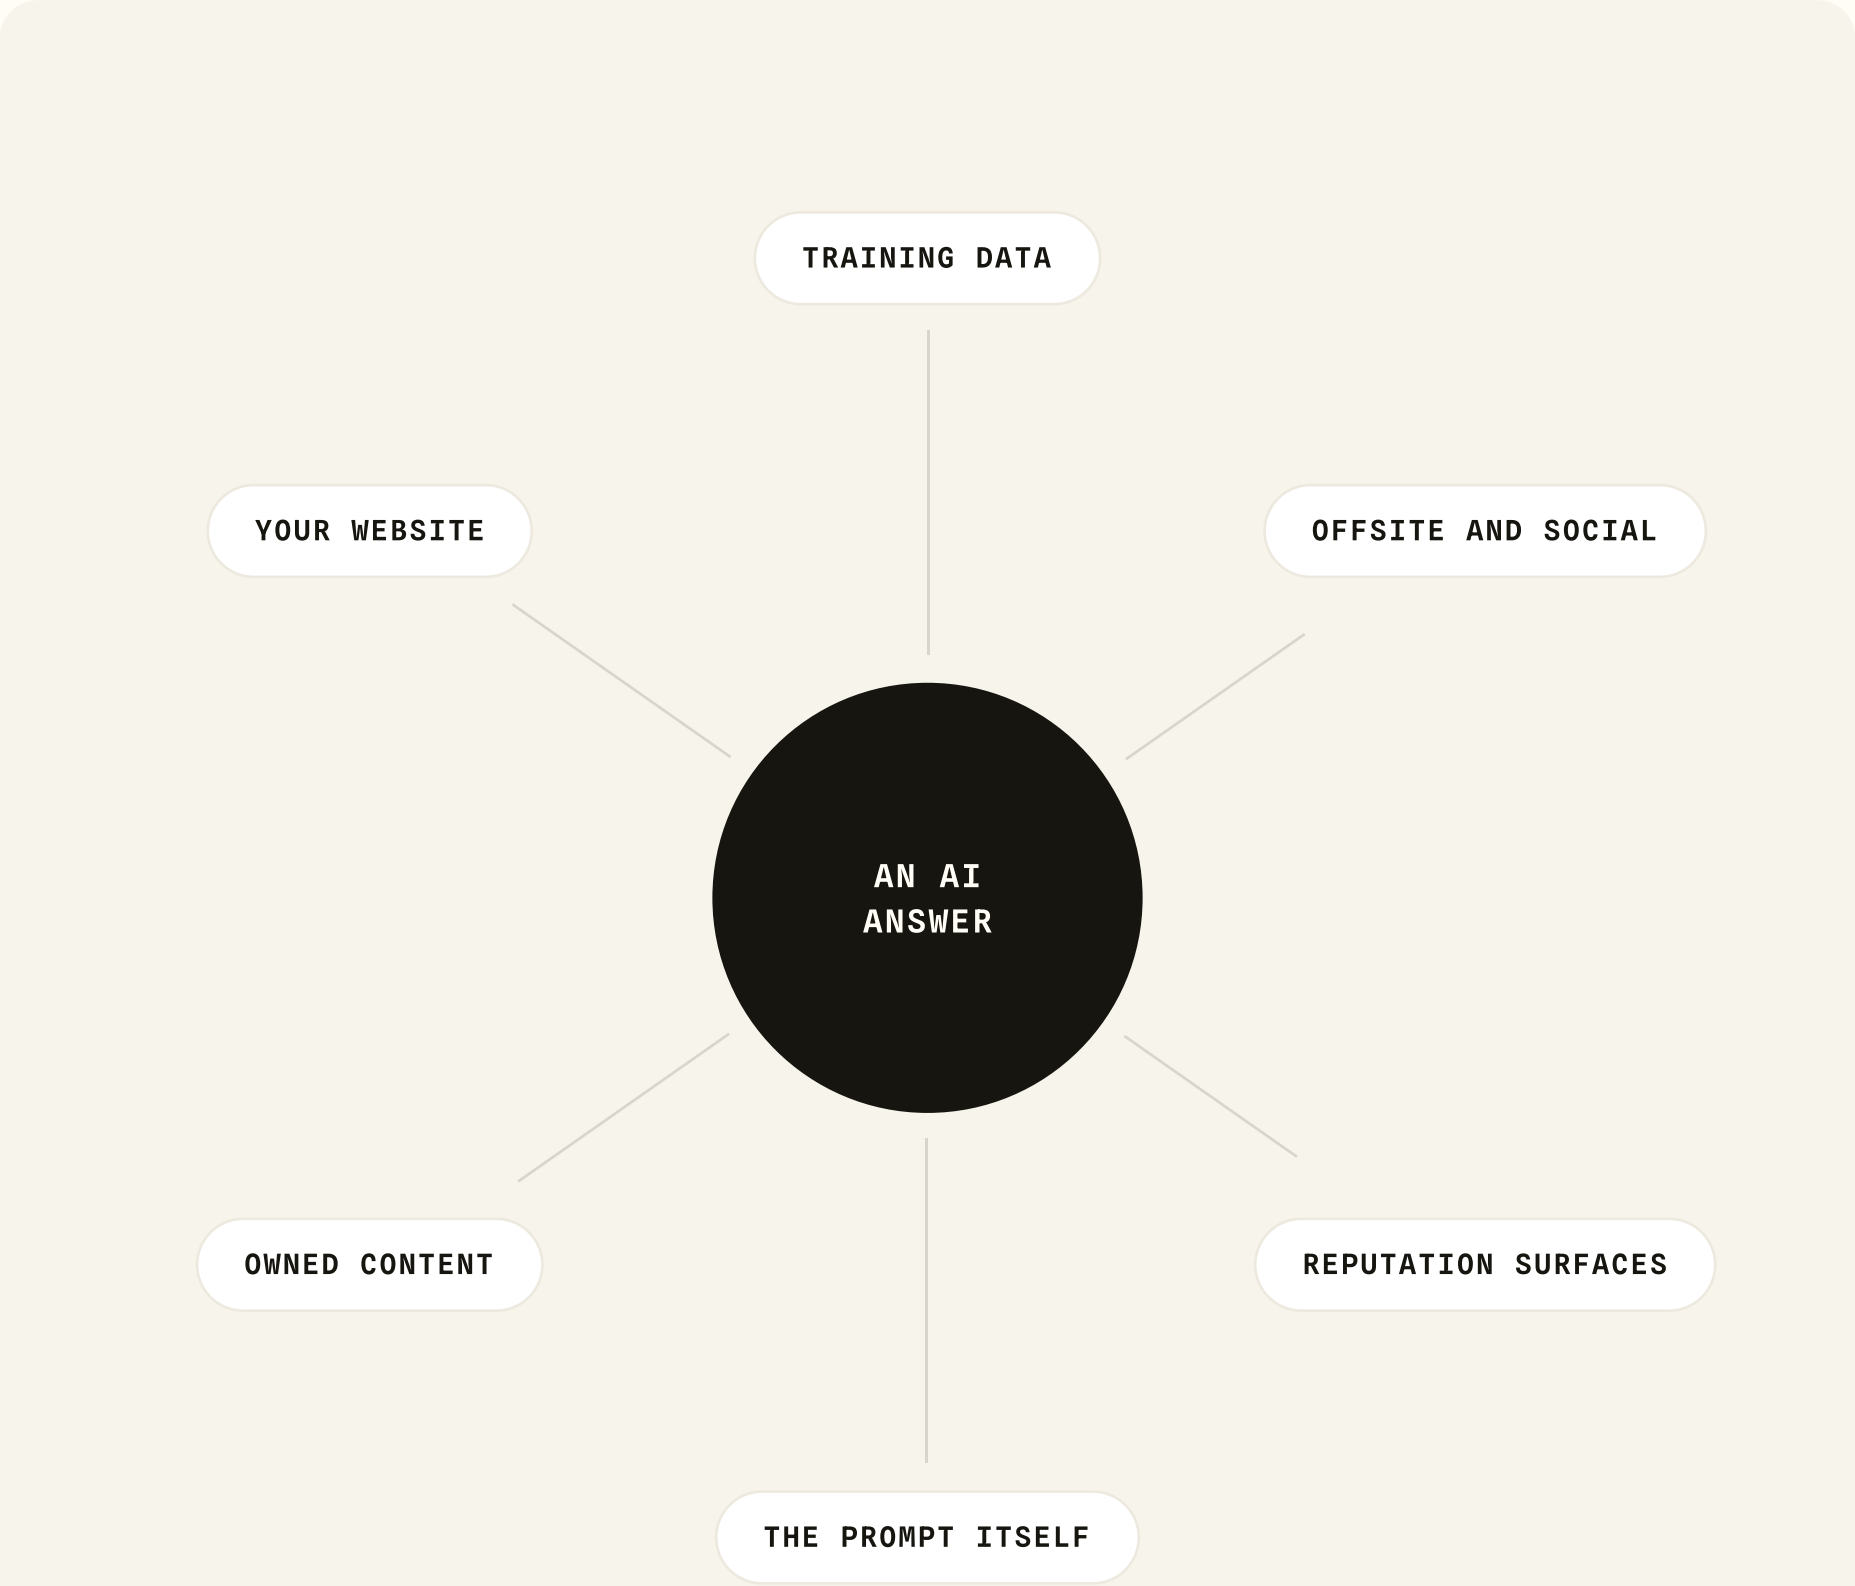
A distinct formation reads as itself from any distance.

06

The Influenceable Web.

Inputs that shape an AI answer.

Models combine several sources into one answer. Brands can influence most of those sources, and Momentic calls that set the Influenceable Web.



The sixth input is the prompt itself. Its sentiment influences answer sentiment, and the entities it treats as central decide which brands the model includes.

Places of corroboration.

On these eight kinds of websites, a third party can confirm a brand's claim. Each confirmation starts with evidence the brand publishes.

01

Editorial coverage

Journalists cite original data and findings.

02

Trade and niche publications

Match specialized prompts more closely than general news does.

03

Professional profiles

The same name, title, and credentials on every profile.

04

Review ecosystems

The review websites your buyers read before choosing.

05

Video

With transcripts and explicit entities.

06

Communities

Accounts from people who have used what you sell.

07

Public registries and listings

Name, address, and category match across directories.

08

Partner and integration ecosystems

Partner pages that confirm specific integrations.

WHERE CITED PAGES ARE HOSTED, ALL CAPTURED CITATIONS



10.7%
BRAND-OWNED

About nine of every ten citations go to websites the brand does not own.

By engine, brand-owned pages received between 5% (Copilot) and 18% (ChatGPT with web search) of citations.

Source: Atlas measurement, 16 brands, 1,288,263 captured citations, 2026-07-13 to 2026-09-13. Owned share by engine: 8 engines that cite sources.



Communities are where first-hand experience genuinely exist. You need to respect that in your strategy.

Tactics Momentic doesn’t use.

Momentic excludes four tactics outright. Momentic builds press on real datasets, qualified experts, and first-hand results instead.

01 / 04

01 **Purchased mentions**
Payment for coverage that is presented to readers as independent.

02 **Staff sockpuppets**
Employees posting as unaffiliated customers or unaffiliated experts.

03 **Seeded forum threads**
Discussion planted by the brand and staged to read as organic.

04 **Listicle insertion**
A ranking position bought inside a third-party roundup.



This exclusion applies to every Momentic engagement.

07

Demand.

Sources for demand research.

A keyword is basically a topic. A prompt adds constraints, context, and the outcome the person wants.

01 / 03

01 **Search demand**

What people already type into a search box, counted in keyword volume and internal site search.

02 **Decision questions**

What a buyer needs answered before acting, taken from sales calls and support tickets.

03 **Fan-out coverage**

The sub-questions a system retrieves separately when it answers one visible prompt.



Validated against keyword data, internal site search, sales calls, and support tickets.

Prompt classifications.

A panel is a fixed, versioned, representative set of prompts per commercial journey, preserved across months.

	PROMPT FAMILY	WHAT IT TESTS
01	**Category discovery**	Unaided inclusion and category authority
02	**Problem and solution**	Ownership of the problem space
03	**Comparison**	Entity clarity and comparative evidence
04	**Constraint**	Attribute-level facts under a budget, integration, or spec
05	**Trust and risk**	Reputation surfaces and balance of evidence
06	**Implementation**	Procedures and operational detail
07	**Local**	Entity data, proximity, reviews
08	**Transactional and agent**	Machine-operable data, action readiness



The questions your customers really ask.

Before comparing, do this.

Wording and denominators decide whether a panel's numbers can be compared across brands or months.

TWO CHECKS BEFORE A PANEL IS COMPARED

Run both checks before any comparison.

01 FRAMING BALANCE

Prompt wording sets sentiment. A panel whose wording nets negative or positive biases every downstream number, so balance the framing to neutral before comparing brands or months.

CHECK

Confirm the panel's wording nets neutral before the first comparison is run.

In Atlas data, prompt type decided caveat rate more than framing did. Category prompts drew caveats about 1% of the time whatever the framing. Branded prompts drew caveats 26% to 59% of the time.

Source: Atlas measurement, 16 brands, 116,532 answers with framing and category labels, 2026-07-13 to 2026-09-14.

02 BRANDED PROMPT SEMANTICS

On a prompt that names the brand, a mention is a given and means nothing. Branded prompts measure whether owned pages get cited and whether answers state caveats. Category prompts measure mention and recommendation. Mixing the two denominators corrupts both.

CHECK

Mention rate and recommendation rate are read only from category prompts.

EACH RATE NAMES ITS OWN DENOMINATOR

owned-citation share

answers citing an owned URL

all answers to branded prompts

mention rate

answers naming the brand

all answers to category prompts

caveat rate

answers stating a caveat

all answers to branded prompts

recommendation rate

answers recommending the brand

all answers to category prompts

Report each rate next to the denominator it was divided by, in the same sentence.

08

Content that survives.

Six kinds of material competitors cannot copy.

If a model can write your article from ten similar pages, it has no reason to cite yours. Publish material competitors cannot copy without doing the same work.

01 **Original data with method**
A dataset you collected, published with the method you used to collect it.

02 **First-hand tests**
A procedure you ran, with the setup and the result.

03 **Real cases with baselines and outcomes**
A named engagement, the starting number, and the ending number.

04 **Expert judgment with named credentials**
A view attributed to a person, with the credential that qualifies it.

05 **Working tools**
Calculators, checkers, or other software a reader can use.

06 **Primary documents**
Filings, transcripts, and specifications reproduced in full.



Primary material, gathered first-hand.

Write sections AI can quote.

A model quotes one section under one heading as a unit, so each section must resolve its own question completely. Models do not select passages by word count, so do not split sections to meet a word target.

01 / 05

01 **Answer immediately**

The direct answer appears in the first sentences under the heading, before context, caveats, or background.

02 **Qualify scope and date**

Label the date, the units, the method, and where the numbers came from, inside the same section.

03 **Evidence the claim**

“Studies show” cannot be checked. Name the source and the year so a model can verify the claim.

04 **Explain why it matters**

State the consequence next to the claim, so the passage still makes sense read away from the page.

05 **Connect to the next question**

The next sub-question gets its own heading, never a new page written to repeat the same answer.

Write each section so it makes sense when a model quotes it alone.

Answer unit

What the 40% study measured.

The 2024 Princeton and KDD paper supplied top search results to the generator first, then changed the content. The 40% figure applies only inside that setup.

01

Supply

The paper gave the generator the top search results before it wrote an answer.

02

Rewrite

One of those supplied sources then had its content changed.

03

Score

The changed source was scored on how visible its attributed text was in the answer.

RELATIVE VISIBILITY CHANGE OF ATTRIBUTED TEXT

What the paper measured, and what it did not.

Attributed text, inside the paper's setup



Conditional on the source already being in the model's context.

WHAT THE PAPER DID NOT MEASURE

Retrieval for a live page

not measured

Traffic

not measured

Revenue

not measured

Do not repeat the claim that adding statistics increases ChatGPT rankings by 40%. It collapses retrieval, context allocation, citation and traffic into one imaginary metric.

Sources: the 2024 Princeton and KDD GEO paper; the 2026 critical survey of 45 studies.

09

Technical eligibility.

Which AI crawlers to allow, and why.

Search inclusion, model training, and user-invoked agent access are three separate decisions, and each system below publishes its own user-agent tokens. Set the policy per decision and per system, and do not copy a generic blocklist.

DECISION	SYSTEM	DOCUMENTED PURPOSE	STANCE
SEARCH	**Googlebot**	Search crawling; generative features retrieve from Search systems	Allow every page you want in Google Search and AI Overviews.
TRAINING	**Google-Extended**	Training and grounding control token; not a crawler, not a ranking input	Decide separately from Search inclusion. Document the decision.
SEARCH	**OAI-SearchBot**	ChatGPT Search discovery	Allow if ChatGPT traffic matters to you, and allow its published IPs through the WAF.
TRAINING	**GPTBot**	Training collection	Decide whether you allow model training, separately from search access.
SEARCH	**PerplexityBot**	Perplexity search and linking, not training	Allow if Perplexity traffic matters to you, and allow its published IPs.
AGENT	**Perplexity-User**	User-requested fetches	User-agent access; documented as generally ignoring robots.txt.
SEARCH AGENT	**ClaudeBot / claude-web**	Anthropic crawling and search-tool fetches	The same purpose split. Verify each token in access logs.

Check crawler access at four layers.

Robots.txt records only the stated policy. A firewall or a CDN rule can refuse a crawler that robots.txt allows, and only the access log's error rate shows that the refusal happened. The audit reads all four layers in order.

01 / 04

01 **Robots.txt**

Written per purpose, because search inclusion, model training, and agent access are separate decisions.

02 **Firewall rules**

Robots.txt allows the bot while the WAF challenges it by IP, and nothing in robots.txt reports the refusal.

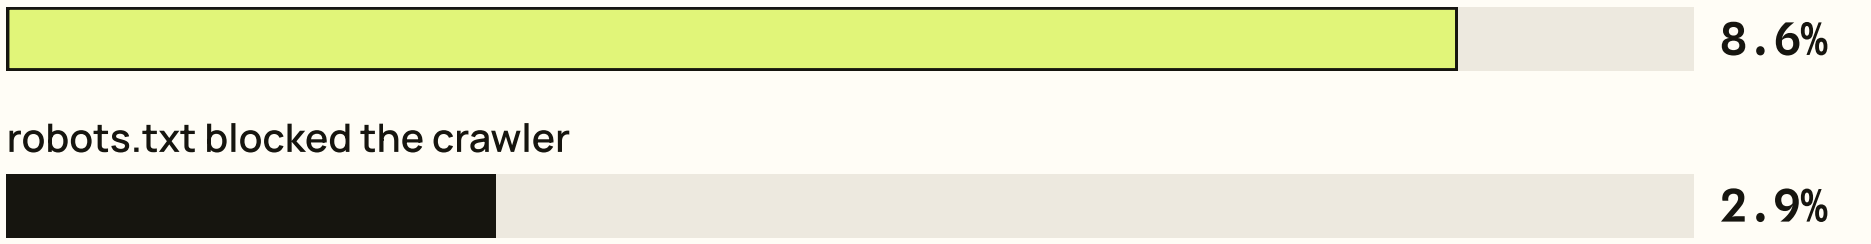
03 **CDN bot management**

Managed bot rules can block what robots.txt permits, so WAF and CDN policy have to match robots intent rather than contradict it.

04 **Access-log error rates**

The only direct measurement: which crawlers came, what they fetched, what they got back, and whether they gave up.

AI CRAWLER CHECKS REFUSED, BY LAYER



Across 98 monitored domains, firewalls and CDNs refused AI crawlers on 8.6% of checks. robots.txt blocked them on 2.9%.

Source: Atlas bot-reachability monitor, 16 brands, 98 domains, latest snapshot, 2026-09-14. Refused = HTTP 403 on a live probe.

Scope: the four layers above, read in the listed order until one refusal is found.

Technical checks.

A crawler must fetch the page, a renderer must see its content, and the index must keep it. If a page fails one of these seven checks, AI-specific work will not fix it.

01 / 07

01 **Stable 200s and clean canonicals**

02 **Important content present in rendered text**

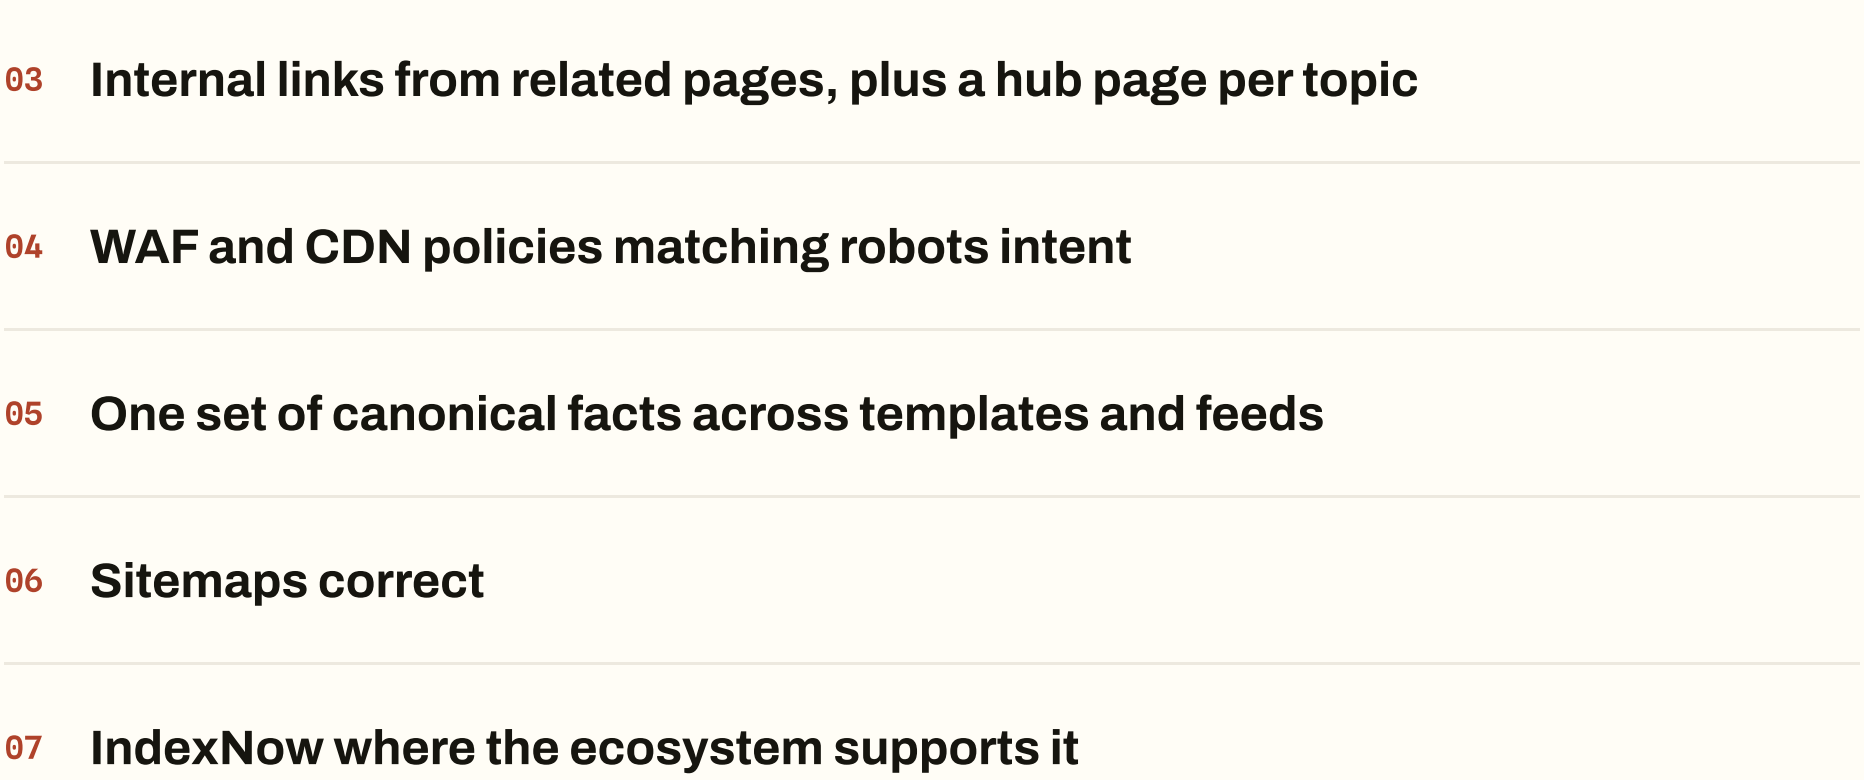
03 **Internal links from related pages, plus a hub page per topic**

04 **WAF and CDN policies matching robots intent**

05 **One set of canonical facts across templates and feeds**

06 **Sitemaps correct**

07 **IndexNow where the ecosystem supports it**



Scope: the seven checks above. A page that fails any one of them is not eligible.

10

Platform punchlists.

Google: AI Overviews and AI Mode.

Google documents how these features retrieve pages, so each check targets a documented behavior. Six checks cover crawl health, subtopic coverage, evidence, structured data, reporting, and the Google-Extended decision.

CONFIRMED

Generative features run on core Search systems, use query fan-out, require normal index and snippet eligibility, and need no special AI schema or AI text file.

SIX CHECKS

01 Ordinary crawl and index health on priority pages

02 Subtopics covered on existing pages, without one page per query

03 Original data or research that other pages do not already contain

04 Structured data valid and matching visible content, for supported features

05 Search Console’s Generative AI performance report reviewed monthly (global since August 2026)

06 Google-Extended decision documented separately from Search

ChatGPT Search.

ChatGPT rewrites a prompt into one or more targeted searches, runs them, and attaches citations to the sources it quotes. Five checks decide whether a page is available to be one of those sources.

CONFIRMED

Inclusion requires OAI-SearchBot access plus its published IPs.

FIVE CHECKS

01 OAI-SearchBot allowed in robots and through the WAF, verified in logs

02 Pages indexed and ranking in regular search

03 Titles, snippets, and page identity precise enough to match the rewritten sub-query

04 GPTBot policy decided separately from search access

05 chatgpt.com referrals tracked in analytics, next to panel visibility for answers that send no click

Perplexity, Claude, Copilot and Bing.

These engines share the same crawl and index foundations. Each enforces its own crawler access and publishes its own first-party report, so the audit checks the three groups separately.

PERPLEXITY

PerplexityBot and its published IPs allowed through the WAF

Current, accurate facts on every page Perplexity may cite

Cited sources and referral visits sampled from Perplexity answers

CLAUDE

Sources built for research tasks: primary documents, method detail, clean provenance

Narrow sub-questions answered precisely, not only the main question

Claude sampled separately; its selection behavior differs by product and configuration

MICROSOFT COPILOT AND BING

Bing Webmaster Tools AI Performance report configured and read monthly

Grounding queries in that report reviewed as a sample of topics

IndexNow live for materially updated pages

Bing Places current for local businesses

11

Commerce, local, agents.

Six product fields AI shopping agents read.

Agents read product data to complete purchases, so identifiers, variants, price, availability, shipping, and returns must match across page, feed, and backend.

SIX FIELDS

01 **Identifiers**

02 **Variants**

03 **Price**

04 **Availability**

05 **Shipping**

06 **Returns**

RULES

01 Product copy organized around the criteria a buyer decides on.

02 Feed errors fixed with the same priority as website errors.

Four facts to get right for local results.

AI engines check a location's name and address, location page, hours, and reviews before they show the business in an answer. Get these four facts right first.



Scope: the four entity facts below, checked on the website and on the brand's listings.

NAP	One name, address, and phone number, identical everywhere it is published.
LOCATION PAGES	Location pages with local details: services, staff, photos, and directions.
HOURS	Hours and exceptions correct and current.
REVIEWS	Reviews monitored, because AI answers quote them.

24 checks so AI agents can use your website.

Browser agents read the DOM the way assistive technology does. The audit scores each item 0 for missing, 1 for partial or inconsistent, and 2 for deliberate, verified and maintained.

8 x 3 = 24

SC

- **SC1** Actions use button or link elements, or a role plus working keyboard handling
- **SC2** Each control's accessible name matches its visible text
- **SC3** Fields carry a programmatic label, an autocomplete token, and the right input type

ST

- **ST1** Disclosures expose aria-expanded; selection exposes aria-current or aria-selected
- **ST2** Loading and busy states reach the tree through aria-busy, disabled, or a live region
- **ST3** The settled tree matches the navigation-time tree for every control in the flow

KB

- **KB1** Every control in the flow is reachable by Tab, in the order it appears on screen
- **KB2** Custom controls respond to Enter and Space
- **KB3** Modals hold focus, return it on close, and close on Escape

VA

- **VA1** Submission failure reaches the tree through role="alert" or aria-live
- **VA2** Invalid fields carry aria-invalid and aria-describedby pointing at the message
- **VA3** The required state in HTML matches the required state in the visible copy

TP

- **TP1** The flow completes with no hover-only, drag-only, or canvas-only step
- **TP2** Steps have stable URLs or stable step names the agent can return to
- **TP3** Consent, chat, and promotion overlays leave the primary control clickable

MT

- **MT1** Price, currency, availability, and stock appear as text, not only in images
- **MT2** Dates, times, time zone, and slot options appear as text or machine-readable values
- **MT3** Cancellation, returns, delivery, and data-handling terms appear at the point of action

CF

- **CF1** A review step shows what will be committed before the irreversible action
- **CF2** The success state is unambiguous and present in the accessibility tree
- **CF3** The failure state is recoverable and keeps what the visitor already typed

HP

- **HP1** Each recommended change was checked against the human experience
- **HP2** No agent-only content: no hidden text, cloaked markup, or agent-only path
- **HP3** Visible labels stay visible for humans after a field is filled

No agent-operability change ships if it makes the page harder for a human to use.

Source: the Momentic agentic-UX audit instrument. Item IDs and wording as published.

12

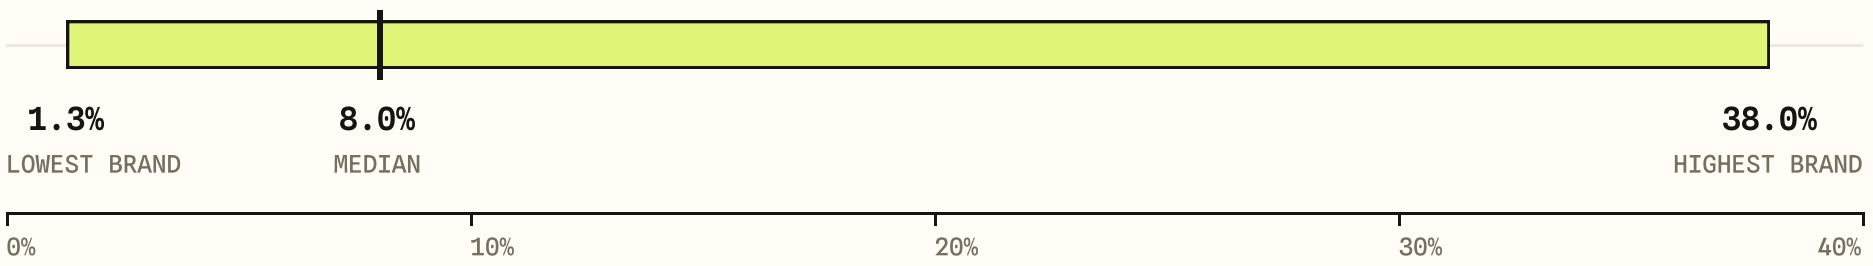
Measurement.

Measure each stage separately.

A composite score does not show which gate failed. Each layer reports one metric, and one named system supplies it.

01	**Reachability** Bot requests, response codes, give-ups	Server, CDN, WAF logs
02	**Google generative** Generative impressions	Search Console
03	**Microsoft citation** Citations, cited pages, grounding queries	Bing Webmaster Tools
04	**Sampled visibility** Mention, citation, recommendation by prompt family	The versioned panel
05	**Representation** Accuracy, sentiment, caveats	Human-reviewed samples
06	**Referral** AI-referred sessions and behavior	Analytics, referral parameters
07	**Demand** Branded search, direct traffic	Search Console plus analytics
08	**Commercial** Leads, revenue, pipeline	CRM and commerce systems

CATEGORY-PROMPT MENTION RATE, EACH OF 16 BRANDS



Source: Atlas measurement, 16 brands (each with at least 30 graded category-prompt answers), 2026-07-13 to 2026-09-13. Range across brands, not a single score.

Four metrics teams often confuse.

Each rate changes independently of the other three, so every report states the denominator next to the number.

01 Mention rate

The answer names the brand.

Denominator: category prompts only. On a branded prompt a mention is a given.

02 Citation rate

The answer uses an owned page as evidence.

Denominator: the prompts run, with branded and category kept apart.

03 Recommendation rate

The answer actively suggests the brand.

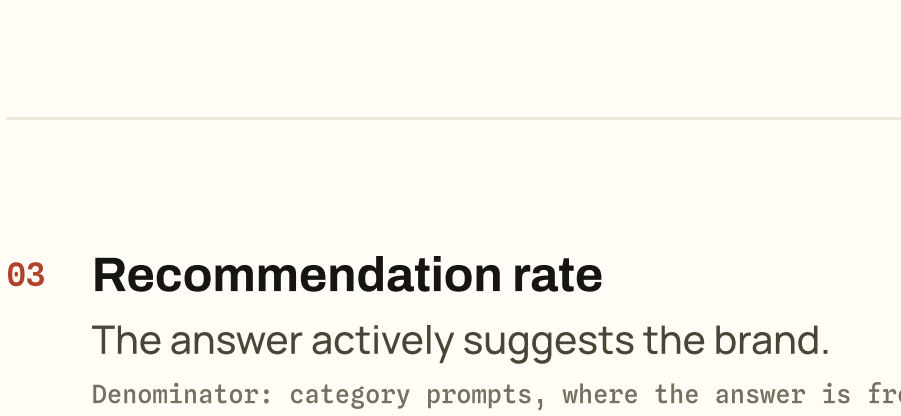
Denominator: category prompts, where the answer is free to choose.

04 Referral rate

A visit arrives from an AI surface.

Denominator: sessions in analytics, never answers.

CATEGORY PROMPTS



A brand was named or cited far more often than it was recommended.

Source: Atlas measurement, 16 brands, 11 engines, 88,574 answers to category prompts, 2026-07-13 to 2026-09-14.

BRANDED PROMPTS

On prompts that name the brand, about three of every four citations go to pages the brand does not own.

Owned-citation share. Source: Atlas measurement, 16 brands, 344,073 citations on branded prompts, 2026-07-13 to 2026-09-13. Caveat rate. Source: Atlas measurement, 16 brands, 64,282 graded branded-prompt answers, 2026-07-13 to 2026-09-13.

Six rules for reporting results.

Check every number against these six rules before publishing it.

01 / 06

01 **Branded prompts do not measure mentions**

Measure owned-citation share and caveat rate on branded prompts. Mention rate belongs to category prompts.

02 **Citations do not attribute to revenue**

Branded search volume and direct visits record influence. The CRM records revenue on its own terms.

03 **Every number names its denominator where it appears**

Compute answer share over every answer to the period's category prompts, never over a brand's own rows.

04 **Framing balances to neutral first**

Balance the wording of a prompt set before running any cross-brand or cross-month comparison.

05 **Two numbers compare only when their run mix matches**

Match engines, product modes, and panel version, or publish a comparability note instead of a trend.

06 **One run does not establish an effect**

Report treated and comparison sets, repeated runs, and the range of the effect.

A number without a stated denominator is not a result.

Reporting register

13

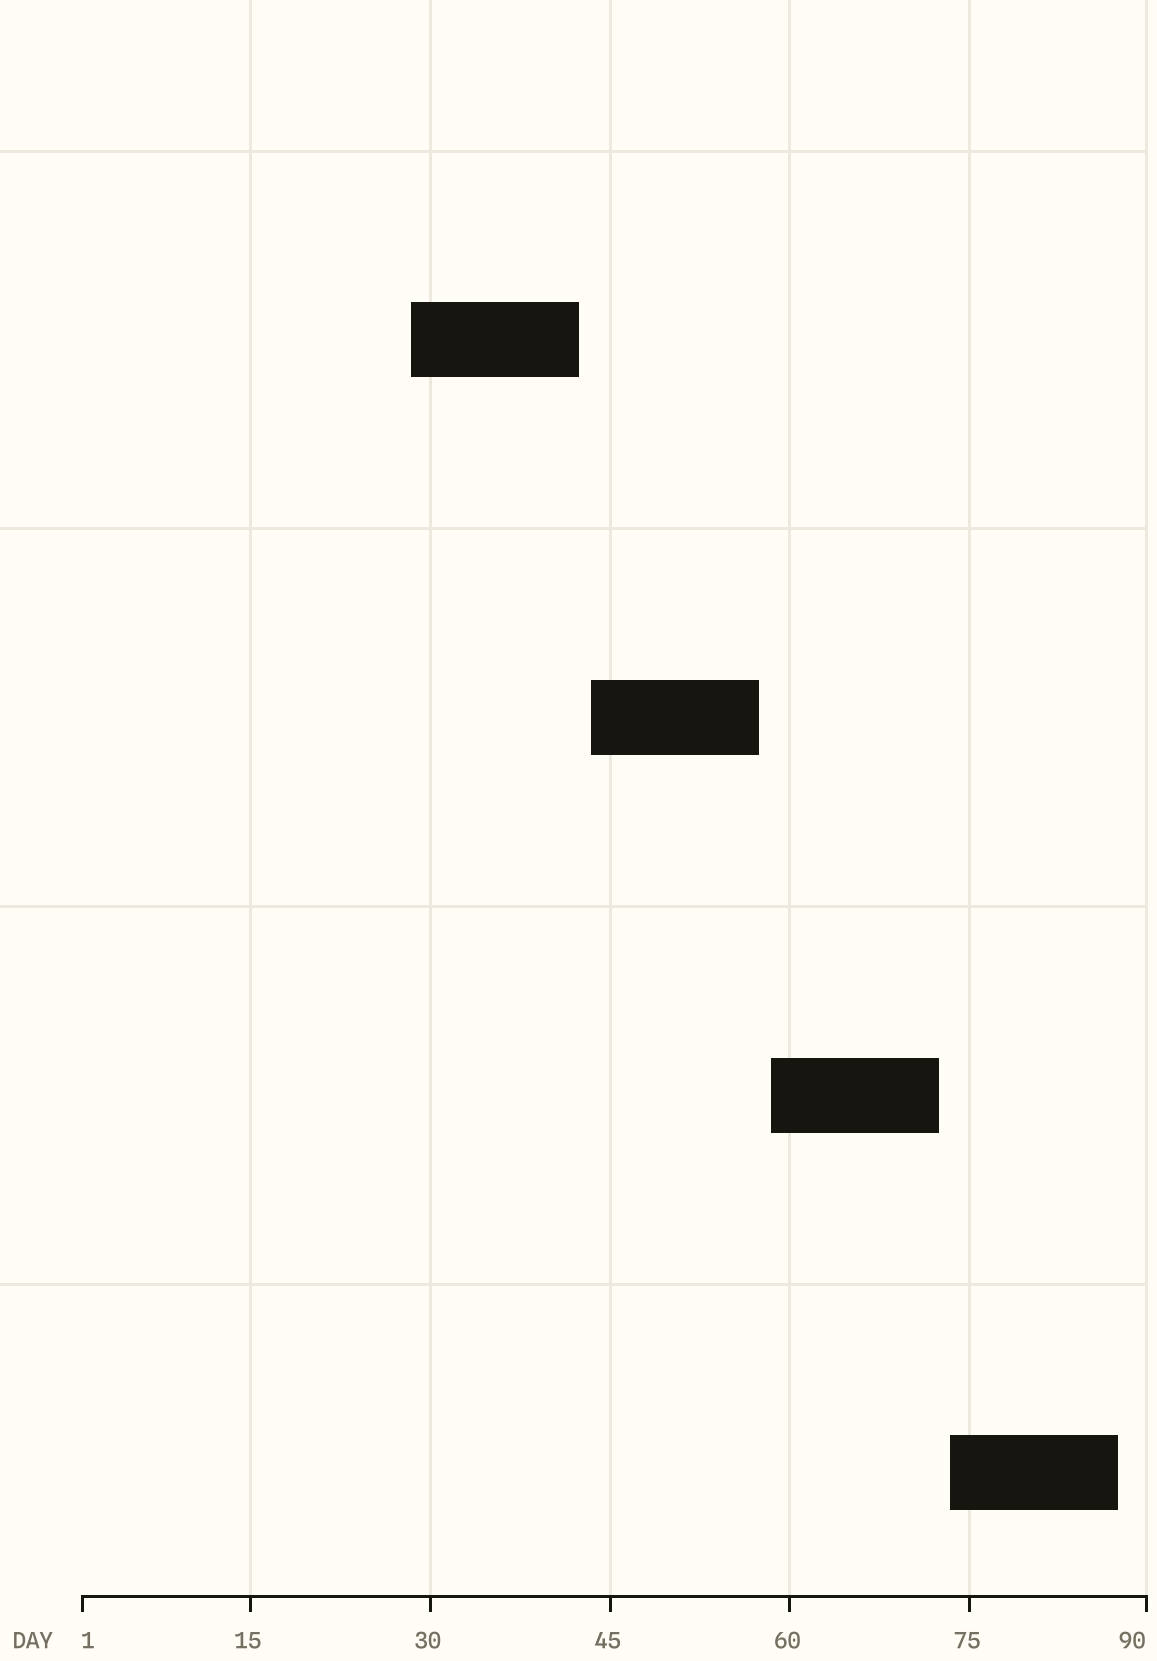
The 90-day track.

What to do in the first 90 days.

Fix crawler access first, then publish original evidence, then run tests. Finish each window's job before starting the next.

15 x 6 = 90

1 to 15	**Baseline and access**	A baseline report of access and visibility
16 to 30	**Retrieval foundation**	Technical fixes plus a URL x Prompt map
31 to 45	**High-leverage pages**	10 to 20 pages materially improved
46 to 60	**Non-commodity assets**	One or two original assets, such as research or a calculator
61 to 75	**Corroboration**	Third-party coverage earned with those assets
76 to 90	**Measure and test**	A log of test results, each with its evidence tier



The Momentic 90-day operating track. Windows drawn at true length on a 90-day axis.

What to start now, test, and skip.

Momentic groups the work into four buckets, from what to start now to what to skip. Each bucket lists five items.

5 x 4 = 20

DO NOW

- Access
- Index health
- Content and reputation distribution
- Page clarity
- Entity consistency

DO NEXT

- Original research
- Evidence-led PR
- Product and local data
- Video with transcripts
- Reference pages updated on a schedule

EXPERIMENT

- Passage variants
- Comparison formats
- Agentic UX
- Selective structured data
- Freshness cadence

SKIP

- Mass generated slop
- Seeded community mentions
- llms.txt as a ranking project
- Schema sold as an AI hack
- Cosmetic date updates



Set priorities before setting dates.

14

The audit.

The ten parts of the audit.

Momentic scores each instrument 0, 1, or 2, records the evidence, and weights each category by business type.

01 / 10

01 **Technical eligibility and rendering**

02 **AI crawler access across all four layers**

03 **Indexation and architecture**

04 **Demand and panel coverage**

05 **Content value and extractability**

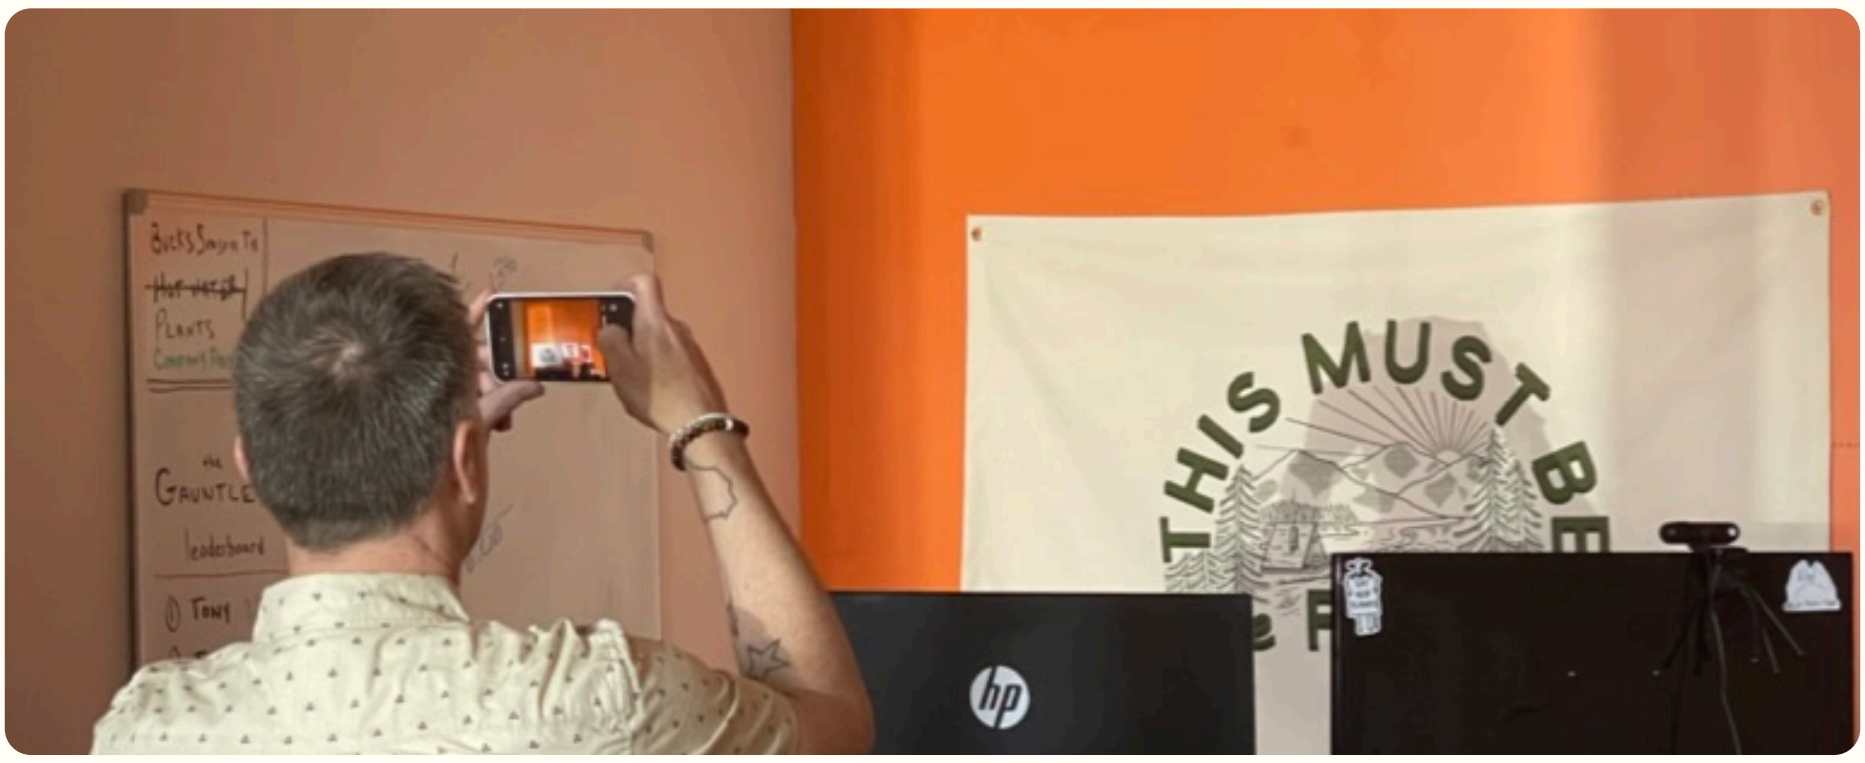
06 **Entity and corroboration**

07 **Structured and feed data**

08 **Platform readiness per engine**

09 **Agentic UX across 24 action-readiness items**

10 **Measurement and governance**



Scoring: 0 missing, 1 partial or inconsistent, 2 deliberate, verified and maintained. Totals are never compared across businesses.

15

The myths.

<footer>ATLAS x MOMENTIC THE AI SEARCH OPERATING MANUAL 55</footer>

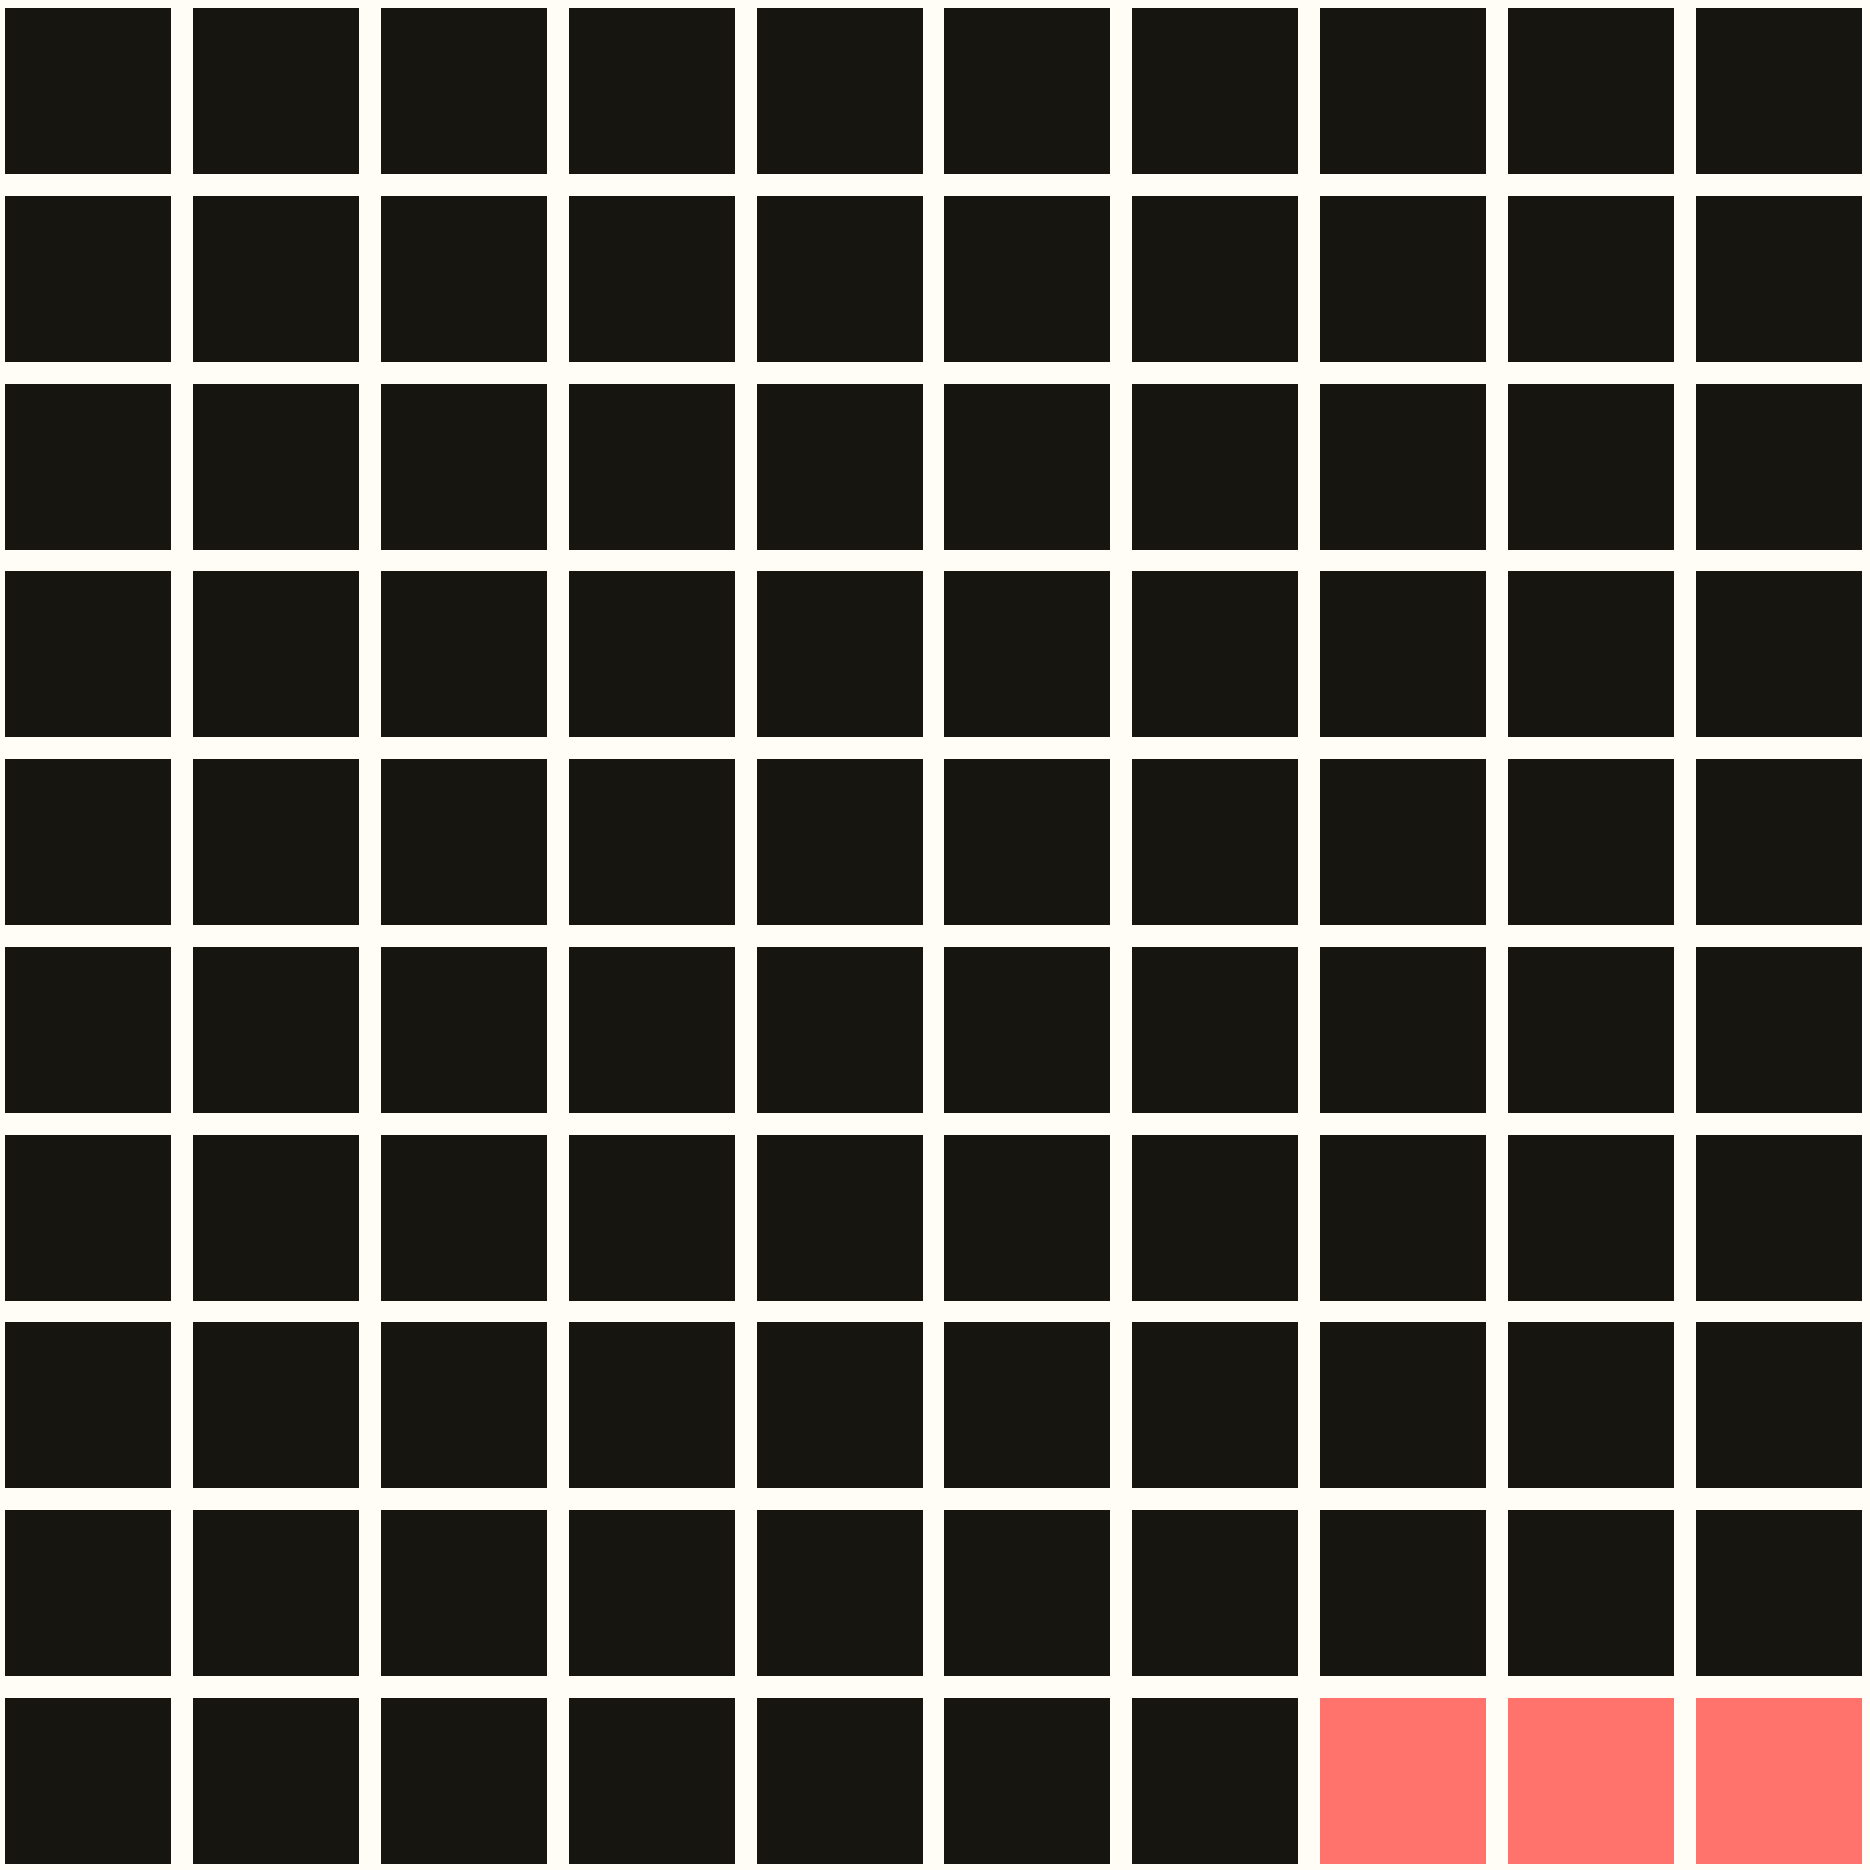
97 of 100 llms.txt files got zero requests.

Ahrefs studied 137,000 sites and logged requests to every valid llms.txt file it found. 97% received no request in the measured period.

97 / 100

EACH SQUARE IS ONE PERCENT OF VALID LLMS.TXT FILES

3 SQUARES = REQUESTED AT LEAST ONCE



Google states that AI text files are unnecessary for its generative features. An llms.txt file can help documentation websites, but no published evidence shows it improves rankings.

Ahrefs llms.txt analysis across 137,000 sites; Google Search Central on AI text files.

Twelve claims to stop repeating.

Each claim below circulates as a tactic. The correction states what published research or Atlas measurement supports.

01 **“Add llms.txt to rank in AI.”**
Google states AI text files are unnecessary for its generative features, and Ahrefs found that 97% of valid files received no requests.

02 **“Schema makes models cite you.”**
Schema is infrastructure for supported features and entity clarity. A tracked study across nearly two thousand pages found citations barely changed after JSON-LD was added.

03 **“Every answer needs a 40-word chunk.”**
Google confirms there is no required length. Write each section to answer one question completely.

04 **“The 40% study proves a ranking boost.”**
The study put the sources into the model’s context first, so its 40% gain does not apply to live retrieval.

05 **“Rank #1 and you will be cited.”**
AI answers regularly cite pages beyond the top search results, and the overlap changes by engine and month.

06 **“Cited means recommended.”**
On category prompts in Atlas data, brands were cited 14.3% of the time and recommended 2.2% of the time. Track each separately.

07 **“Fresh always wins.”**
Freshness matters for facts that change, such as prices. Evidence decides stable facts. Changing only the date does nothing.

08 **“Reddit is the ranking factor.”**
Real community posts help on some prompts and engines. Fake posts break platform rules and get removed.

09 **“Traffic no longer matters.”**
Referral counts miss answers that send no click, but the visits that arrive are real. Track answer metrics alongside traffic and revenue.

10 **“AI search needs its own content team.”**
The work spans search, content, PR, product data, and analytics. A separate team tends to publish duplicate pages with conflicting facts.

11 **“Write for machines now.”**
Writing that is clear to people is also what models extract best. Google asks for content written for people first.

12 **“There is one algorithm to reverse-engineer.”**
Each product uses different models, retrieval systems, and modes. Measure each one repeatedly and report ranges.

16

The one-page version.

The manual in ten steps.

Run the whole manual in this order.

01 / 10

01 **Choose the task**
Pick a task a buyer needs done.

02 **Check reachability**
All four access layers, verified in logs.

03 **Map the fan-out**
Sub-questions, entities, criteria, evidence.

04 **Own a unique fact**
Publish what no model can synthesize from competitor pages.

05 **Make it explicit**
Answer first, then label dates, units, methods, provenance.

06 **Earn corroboration**
Give journalists, partners, and reviewers evidence they can cite.

07 **Run the panel**
Same prompts, every engine, framed neutral: mention, citation, recommendation, accuracy.

08 **Read first-party signals**
Search Console generative, Bing AI Performance, server logs, referrals.

09 **Test with controls**
Comparison sets, repeated runs, effect ranges.

10 **Optimize the action**
Make the next step, such as booking or buying, easy for a person or an agent.

ATLAS

x Momentic

Be the source behind the answer.

AI search is still search: a contest to be the most useful, trusted, retrievable source for a real human need. The interface changed. The standard went up.