

CONTENT & SEARCH · A FIELD GUIDE

The AI Visibility Framework

Eight checks that shape whether AI answers cite your page, or someone else's.

AI ANSWER

4 SOURCES

What's the best scheduling software for a small clinic?

[Product] is scheduling software for clinics with 5 to 50 staff, with setup in a few days. ¹

1

yourbrand.com/pricing

2

competitor-a.com/blog

3

review-site.com/best-tools

CITED

GOOGLE AI OVERVIEWS · AI MODE · CHATGPT · PERPLEXITY · GEMINI



Marina Comisi
Founder, Paxton Studio


BEFORE YOU START

Your buyers ask AI for a shortlist. Is your name *on it?*

A buyer asks ChatGPT for the best [your category] for a company like theirs. The answer names three competitors and links to their pages. You're not on the shortlist, and nobody on your team knows it happened. If your buyers research this way, that conversation takes place before the first sales call.

Part of the reason is mechanical. In 18,012 verified ChatGPT citations, 44.2% pointed to the first 30% of the page. AI systems lift the passage that answers best, so a page can rank on Google and never be quoted.

[Kevin Indig, Growth Memo · Feb. 2026 ↗](#)



WHAT'S INSIDE

- 01

Eight levers, ordered by impact and control
- 02

Three before-and-after rewrites to copy
- 03

A scorecard for your ten key pages
- 04

A tracking method, 20 prompts, a tracker

HOW TO READ THE NUMBERS

Every figure carries its source and date. Most are correlations, measured on one engine at one point in time, and the models change faster than the studies. We flag that each time, so you can defend this framework in front of a CFO.

CONTENTS

Is this for you?	03	04 · The scorecard	13
01 · How an AI answer picks its sources	04	05 · The order we'd work in	14
02 · The eight levers	05	06 · What to track, and revenue	15
03 · Three rewrites	11	Glossary and 20 prompt templates	18
Our live pilot: week 0	12	Tracker template and sources	19

BEFORE YOU SPEND A BUDGET ON IT

Is this the right problem *to fix now?*

This work pays off when buyers ask questions before they talk to anyone. It can wait when the problem sits somewhere else.

START HERE IF

- ✓ Your buyers compare options before contacting sales
They ask "which is best for...", check prices, read comparisons.
- ✓ Competitors already appear in AI answers for your category
The shortlist is forming, with or without you.
- ✓ You publish pages that answer buyer questions
Pricing, comparisons, guides, FAQs.
- ✓ A few pages bring most of your commercial traffic
Ten pages done properly will move the needle.

START ELSEWHERE IF

- Your offer or positioning isn't clear yet
AI answers will repeat the confusion. Fix the message first.
- Your site doesn't convert the traffic it already gets
More visibility on a leaky page is wasted.
- Most buyers come from referrals, events or your network
Invest where they actually decide.
- Nobody can update pages for the next two months
Plan the work before you measure it.

A 30-MINUTE WAY TO FIND OUT

Ask ChatGPT and Perplexity five questions your buyers ask, without your brand name. Note who gets named. If competitors show up and you don't, this guide is for you.

5 QUESTIONS × 2 ENGINES


10 answers

Starter path, page 15

01 · THE MECHANISM

How an AI answer picks *its sources*


Most brands lose at the passage, after their page has already been found. Most engines follow four broad moves.



STEP 1

The query fans out

One question becomes several related searches. You compete on questions nobody typed.



STEP 2

Pages are retrieved

From each engine's index or crawlers. A page that can't be read never enters.



STEP 3 · WHERE MOST PAGES FAIL

Passages are selected

The model scores chunks of text and keeps the ones that answer cleanly.



STEP 4

A source is cited

Usually the page that supplied the passage. Sometimes one that only confirmed it.

[Google Search Central · AI features and your website ↗](#)

Picture a research assistant on deadline, with ten printouts and *thirty seconds*.

They don't read the pages. They scan each one for the paragraph that answers the question, check it has a name or a number they can quote, and staple it to their reply. This guide is about being the paragraph that gets stapled.

WHAT YOU CAN SKIP

Google states there are no special optimizations, no AI text files and no special schema needed for AI Overviews or AI Mode. Indexed and eligible for a snippet is the bar.

[Source · Google ↗](#)

WHAT THE STUDIES CAN'T TELL YOU

The research here is one 2023 experiment, vendor correlation studies and platform docs. None proves a change moves citations by a set percentage. Read each number as a direction, then test it.

Testing method · page 15

02 · THE EIGHT LEVERS

Eight levers, in the order *we'd pull them.*

Lever 0 comes first: nothing else counts if the crawler can't read the page. Levers 1 to 3 carry the strongest evidence, lever 4 is the core editorial priority, and all four sit in your hands.

#	LEVER	WHAT THE EVIDENCE SAYS	PRIORITY
0	Readable by AI crawlers	OpenAI, Anthropic and Perplexity crawlers don't execute JavaScript	PREREQUISITE
1	The answer sits high	44.2% of ChatGPT citations come from the first 30% of the page	STRONGEST EVIDENCE
2	Attributed quotations	Best method in the GEO experiment: about +41% visibility	STRONGEST EVIDENCE
3	Specific, sourced numbers	Statistics +31%, cited sources +27% in the same experiment	STRONGEST EVIDENCE
4	Information nobody else has	Rewritten consensus gives the model no reason to pick you	CORE EDITORIAL PRIORITY
5	Structure a model can lift	Cited text is about twice as likely to sit under a question	MULTIPLIER
6	Entity clarity and mentions	Brand mentions correlate about 3x more than backlink count	SLOW, COMPOUNDING
7	Freshness where topics move	ChatGPT cites pages about 458 days newer than Google organic	ENGINE-DEPENDENT

LEVER 0 OF 8 PREREQUISITE

Make sure AI crawlers *can read* you

A page whose text only appears after scripts run is blank to most AI crawlers. Gemini is the exception: it uses Googlebot, which renders the page.

WHAT A BROWSER SHOWS

WHAT GPTBOT, CLAUDEBOT OR PERPLEXITYBOT SEE

<DIV ID="ROOT"></DIV>

Nothing to quote

when the text loads through JavaScript

CRAWLER	USED BY	RUNS JAVASCRIPT
GPTBot · OAI-SearchBot · ChatGPT-User	OpenAI	No
ClaudeBot	Anthropic	No
PerplexityBot	Perplexity	No
Googlebot	Google, Gemini	Yes

WORTH KNOWING

~35%

of ChatGPT and Claude crawler fetches hit 404 pages in Vercel's data. Old redirects cost you reads.

[Vercel, The rise of the AI crawler · Dec. 2024 ↗](#)

THE CHECK

10 MIN · ONCE PER SITE

☐ Load your top page with JavaScript off, or run `curl`. Is the main answer in the raw HTML?

☐ Open robots.txt and your CDN's bot settings. Are the search crawlers you want allowed?


☐ List broken URLs and old redirects, and fix the ones your key pages link to.

LEVER 1 OF 8


STRONGEST EVIDENCE

Put the answer where the model *looks first*

WHERE ON THE PAGE CITATIONS COME FROM



INSIDE A PARAGRAPH



18,012 verified ChatGPT citations. Correlation, one engine.

[Kevin Indig, Growth Memo · Feb. 2026 ↗](#)

Nearly half of ChatGPT's citations point to the top third of a page. Inside a paragraph, the pattern flips.

So the page's main answer belongs near the top, and each section should open with a short paragraph that answers its heading in full. It doesn't need to be a one-line slogan: the model lifts whole sentences from the middle of a paragraph more often than first lines.

Cited passages were also nearly twice as likely to contain a plain definition, the "X is..." sentence most teams cut as too obvious.

THE CHECK

5 MIN · PER PAGE

- ☐ Open your ten most important pages. Read only the first paragraph under each H2.
- ☐ Ask: would a buyer get the answer without scrolling? If not, move it up.
- ☐ Score it on the scorecard, page 13: 0, 1 or 2.

BEFORE


Some thoughts on pricing. Pricing is one of the most debated topics in our industry, and there's no one-size-fits-all answer...

AFTER

How much does a brand identity cost in 2026? For an early-stage startup, it typically costs [range], depending on [two named factors] ([source, date]).

LEVERS 2 AND 3 OF 8 **STRONGEST EVIDENCE**

Give the model a quote and a number *it can lift*



LEVER 2 · QUOTATIONS WITH A NAME ON THEM

Adding relevant quotes was the best of nine methods. Use real ones: a named expert, a customer with their role, or a statement your founder has already published. An invented quote is the first thing a skeptical reader checks.

☐ One quote per key page, with full name, role and context.

LEVER 3 · NUMBERS INSTEAD OF ADJECTIVES

Statistics and cited sources come next. Keyword stuffing does worse than nothing. Heavily cited text also averaged 20.6% proper nouns, against 5 to 8% in ordinary English: models favor text full of named things.

☐ Count sourced, dated numbers. Then count adjectives doing a number's job.

LEVER 4 OF 8 CORE EDITORIAL PRIORITY

Bring what the other nine pages *don't have*



Ten pages say the same thing: the model can cite any of them. Give it a reason to pick yours.

- 01 Data from your own work
- 02 A test you ran, with its result
- 03 A process detail only a practitioner knows
- 04 A position against common advice, and why

This is the lever no study can measure for you, and the one competitors can't copy by *rewriting your page*.

LEVER 5 OF 8 MULTIPLIER

Shape pages so a passage *lifts out cleanly*

2x

more likely to contain a question mark, in cited content

78.4%

of question-linked citations came from headings

Descriptive headings phrased the way buyers ask, short paragraphs, numbered steps, tables and FAQ blocks give the model clean units to extract. Schema is hygiene: Google says no special markup is required for its AI features.

A PAGE A MODEL CAN LIFT

How much does it cost?

H2 AS A QUESTION

ANSWER FIRST


A TABLE

NUMBERED STEPS

+ Is it worth it for...? FAQ

LEVER 6 OF 8 SLOW, COMPOUNDING

Be the same name everywhere, and get *mentioned off-site*



[Ahrefs, AI Overview brand visibility factors · May 2025 ↗](#)

Use the same brand, product and founder names on your site, LinkedIn, directories, review platforms and press. Then earn mentions on the pages your buyers already read.

WHERE TO START
Your tracker lists the third-party pages the engines already cite for your prompts. Get onto five of them: you know the engines read them.



03 · BEFORE AND AFTER

Three rewrites *you can copy*

The brands and figures are placeholders. Brackets mark what you replace with your own verified data; never fill them with estimates. The **highlight** marks what changed.

01 · A SAAS HOMEPAGE SECTION LEVER 1 LEVER 3

BEFORE

Why choose us. Our platform delivers a seamless, best-in-class experience trusted by leading teams worldwide.

AFTER

Who is [Product] for? [Product] is scheduling software for clinics with 5 to 50 staff. Setup takes [X] days and pricing starts at [\$X] per seat per month (as of [month, year]).

The heading becomes a question a buyer types. The first sentence defines, the second carries two numbers to quote.

02 · A SERVICE PAGE LEVER 2 LEVER 3 LEVER 4

BEFORE

We help ambitious brands grow with data-driven strategies tailored to their goals.

AFTER

Across our last [N] projects, the median time from signed brief to live landing page was [X] days. "[Exact quote about the result]," says [Full name], [role] at [company].

One first-party number, one attributed quotation. Neither exists anywhere else on the web.

03 · A BLOG ARTICLE LEVER 1 LEVER 5

BEFORE

Some thoughts on pricing. Pricing is one of the most debated topics in our industry, and there's no one-size-fits-all answer. Let's look at the factors...

AFTER

How much does a brand identity cost in 2026? For an early-stage startup, it typically costs [range], depending on [two named factors] ([source, date]). The table below breaks it down by scope.

The warm-up goes. The answer moves into the first paragraph, and a table carries the detail.

OUR LIVE PILOT · WEEK 0

We're testing the framework *on ourselves*

Paxton Studio relaunched with a new positioning this month, so we're running a 12-week pilot on our own site. This is week 0, published before any result. The six questions below are a first observation, not the full 180-answer protocol.

OCT. 20, 2026

Week 0

On 6 unbranded buyer questions, ChatGPT shortlisted other firms. The pilot will test whether these changes improve our inclusion.

CHATGPT · "WHAT IS PAXTON STUDIO?"

Led with a namesake photography studio, then described a Paris design agency at the wrong domain, paxtonstudio.com IDENTITY CONFUSION

PERPLEXITY · "WHAT IS PAXTON STUDIO?"

An accurate description of today's offer. Only the page title still carried the pre-relaunch tagline OLD TITLE

WHAT WE'RE CHANGING

MOSTLY LEVERS 6 AND 1

- 1

Five offer pages, answers first
- 2

Titles and descriptions on every page
- 3

One name, domain and line on every profile
- 4

Outreach to the comparison pages cited instead of us



TARGETS — NOT RESULTS

SET BEFORE THE PILOT STARTED

+10 pts	35 of 36	0
unbranded mention rate vs the average of the two baseline readings, held at weeks 8 and 12	branded answers describe us correctly	major invented facts: prices, clients or services

We'll publish the week-12 reading in the next edition, *whatever it shows*.

04 · PRINT THIS PAGE

Grade a page *out of 16*

One sheet per page. Pick your 10 most important pages (traffic plus buyer intent), score each lever 0, 1 or 2. About five minutes a page.

PAGE URL

DATE

#	LEVER	0	1	2	SCORE
0	Readable by AI crawlers	Main text missing from raw HTML, or crawler blocked	Text present, some key content behind scripts or tabs	Full text in raw HTML, search crawlers allowed	
1	Answer sits high	Main answer below the first screen	Answer high, sections open with a warm-up	Page and every section open with a complete answer	
2	Attributed quotations	None	Quote without a full name or role	One quote or more with name, role and context	
3	Sourced numbers	No numbers, or numbers with no source	1 to 2 sourced numbers	3+ numbers, each with a source and a date	
4	Original information	Says what competitors say	One original element (opinion, example)	First-party data, a test or a process detail	
5	Liftable structure	Long blocks, vague headings	Some lists or tables, headings partly descriptive	Question headings, short paragraphs, a table or FAQ	
6	Entity clarity and mentions	Names inconsistent, no third-party mentions	Consistent names, few mentions	Consistent names, regular third-party mentions	
7	Freshness	Time-sensitive page, not updated in 12+ months	Updated, date not visible	Real updates, date visible (or topic doesn't age)	
TOTAL / 16					

13–16

MAINTAIN

Keep it fresh. Put the effort into lever 6 and tracking.

8–12

FIX NEXT

Fix the 0s, starting with levers 1 to 4.

0–7

FIX FIRST

Top of the list. A rewrite usually beats a patch.

0 on lever 0

Stop. Fix rendering or access before anything else.

The thresholds are our working rule of thumb, not a statistical cut-off. Their job is to rank pages against each other so you know where to start.

05 · A FIRST CYCLE, SIX TO EIGHT WEEKS

Ten pages, done properly, *in eight weeks*

STEP	W1	W2	W3	W4	W5	W6	W7	W8
Pick 10–15 pages and the prompts, run the baseline MARKETING LEAD · ½ DAY								
Clear lever 0 (dev ticket, day one) DEVELOPER · 1–3 DAYS								
Rewrite for levers 1 to 4 WRITER + EXPERT · 2–4 H A PAGE								
Lever 6, in parallel MARKETING OR FOUNDER · 2–3 H A WEEK								
Levers 5 and 7 on the same pages WRITER, DEV FOR TEMPLATES · ~1 H A PAGE								
Re-run the prompts, compare page by page MARKETING LEAD · ½ DAY								

Depth on ten pages beats a light pass on *two hundred*.

- 1

Write the prompt set and run the first baseline before you change a word, or you'll have nothing to compare against.
- 2

Rendering, robots.txt, CDN bot rules, broken URLs. Often a developer ticket: open it on day one.
- 3

Answers up top, quotations, sourced numbers, one original element per page.
- 4

Mentions take months. Align names across your profiles, then pitch the third-party pages your buyers read.
- 5

Headings, tables, FAQ and update dates go on the same pages, once the substance is right.
- 6

Keep what moved, revisit what didn't, then add the next ten pages.

06 · MEASUREMENT

Citations are the number *to watch*

Traffic from AI answers is still poorly attributed, so citations and share of voice are your leading indicators. AI answers change from one run to the next: a single screenshot proves nothing.

YOUR PROMPT SET · A METHOD THAT HOLDS UP

20–40

prompts, written the way buyers ask



4

engines, every cycle



3

runs per prompt and engine, logged out, same country, same day

4 wks

between cycles. Log every run in the tracker, page 19.

$$\text{Citation rate} = \frac{\text{runs that cite you}}{\text{total runs}}$$

Report it per engine and per prompt group, next to the three competitors that appear most often instead of you.

START LIGHT

The full protocol means 240 to 480 answers to read per cycle. Start with 10 prompts × 2 engines (ChatGPT, Perplexity) × 1 run: 20 answers, about an hour. If competitors show up and you don't, move to the full protocol.

A spreadsheet is enough to start, and it makes you read the answers. The third-party pages that keep showing up become your outreach list for lever 6.

GOOGLE SEARCH CONSOLE

Generative AI performance reports

Rolled out to all sites worldwide since Aug. 31, 2026. Impressions in AI Overviews and AI Mode, by page, country and device.

No report? The site may have too few impressions, or be excluded from generative AI features. Your prompt set still carries the measurement.

[Search Console Help](#) · checked Oct. 20, 2026



06 · FROM CITATIONS TO PIPELINE

Follow the citation all the way *to the pipeline*

Citations are a leading indicator. Three signals connect them to the pipeline, none of them perfect on its own.

01 · VISITS

AI referrals

Sessions and leads from chatgpt.com, perplexity.ai, gemini.google.com and copilot.microsoft.com in GA4.

02 · LEADS

Self-reported

Add "ChatGPT or another AI assistant" to the "How did you hear about us?" field, and ask on first calls.

03 · DEMAND

Branded search

Branded impressions in Search Console often rise after being named in answers. Slow, but free.

A CAVEAT, OCTOBER 2026

Practitioners report that ChatGPT links recently stopped passing a referrer and their UTM tag, so part of that traffic now shows as Direct in analytics. OpenAI hasn't confirmed it. Watch Direct traffic to your key pages alongside the referral report.
[Relevant Audience · Oct. 2026 ↗](#)


THE MONTHLY LINE FOR LEADERSHIP		TEMPLATE	
METRIC		LAST MONTH	THIS MONTH
Citation rate, unbranded questions		[x]%	[x]%
Named in comparison answers		[x] of [n]	[x] of [n]
Branded answers that are accurate		[x] of [n]	[x] of [n]
AI referral sessions (GA4)		[x]	[x]
Leads saying "found you via AI"		[x]	[x]
Branded search impressions		[x]	[x]

The link will stay imperfect. Three weak signals pointing the same way beat *one strong claim*.

If you do one thing *this week*

Open your top page with JavaScript off, then read its first paragraph under every heading. That half hour tells you whether the problem is access, placement or substance, and each of those has a different fix.

The studies will keep changing, and so will the engines. The habits in this guide have held across every dataset we've read so far: say the answer early, back it with something checkable, and give the model a reason to pick you over the nine other pages saying the same thing.




ABOUT THE AUTHOR

Marina Comisi

Founder of Paxton Studio, a marketing strategy studio that also executes. For more than eight years she has worked across strategy, growth, websites, tracking and content, for companies from first launches to global brands.

FROM FIRST LAUNCHES TO GLOBAL BRANDS



Want a second pair of eyes *on your pages?*

Send us three key pages. We'll score them, name the main gaps and send back a prioritized list of fixes.

hello@paxton.studio

Pāxton *studio*®

paxton.studio ↗

WORDS USED IN THIS GUIDE

Citation	A link or source name shown inside an AI answer
Passage	The chunk of your page the model uses: a paragraph, a list, a table row
Crawler	The robot that downloads your page for an engine
Raw HTML	Your page as the server sends it, before any script runs

Entity	A thing the model can name: your brand, a product, a person
Information gain	What your page adds that the other top pages don't
Share of voice	How often you're cited against competitors, on the same prompts
Schema	Code that labels content for machines (author, price, FAQ)

APPENDIX A

20 prompt templates *to start with*

Replace the brackets with your category, buyer and market. Keep the wording your buyers use in sales calls and reviews, not your internal vocabulary.

DISCOVERY · UNBRANDED

- 01 What is the best [category] for [buyer type]?
- 02 Which [category] do [buyer type] in [country or city] use?
- 03 What should I look for in a [category]?
- 04 Is [category] worth it for a company of [size]?
- 05 How do I choose a [category] when I have [constraint]?

PROBLEM · UNBRANDED

- 06 How do I fix [problem your product solves]?
- 07 Why is [symptom] happening, and what are the options?
- 08 What's the fastest way to [job to be done]?
- 09 How do other [buyer type] handle [problem]?
- 10 What tools help with [problem]?

COMPARISON AND PRICE · UNBRANDED

- 11 [Competitor A] vs [Competitor B]: which is better for [use case]?
- 12 Alternatives to [market leader] for [buyer type]
- 13 How much does [category] cost in [year]?
- 14 Cheapest [category] that still does [key feature]
- 15 [Category] with [specific feature or integration]

BRANDED

- 16 What is [Brand]?
- 17 Is [Brand] good for [use case]?
- 18 [Brand] vs [main competitor]
- 19 [Brand] pricing
- 20 [Brand] reviews

APPENDIX B

The tracker, *one row per run*

Copy it into a spreadsheet. Citation rate = runs where your brand is cited ÷ total runs, per engine and per prompt group.

DATE	#	ENGINE	RUN	CITED	YOUR PAGE	POS.	COMPETITORS CITED
2026-11-02	1	ChatGPT	1	No	—	—	[Competitor A], [Competitor B]
2026-11-02	1	Perplexity	1	Yes	/pricing	3	[Competitor A]

SOURCES · ALL OPENED AND CHECKED ON OCTOBER 20, 2026

Aggarwal, Murahari, Rajpurohit, Kalyan, Narasimhan, Deshpande — GEO: Generative Engine Optimization ↗ . Percentages computed from Table 1.	KDD 2024
Kevin Indig, Growth Memo — citation position study, via Search Engine Land ↗	FEB. 2026
Ahrefs — AI Overview brand visibility factors, 75K brands ↗	MAY 2025
Ahrefs — AI assistants prefer to cite fresher content ↗	JULY 2025
Vercel — The rise of the AI crawler ↗	DEC. 2024
Google Search Central — AI features and your website ↗	DOCUMENTATION
Google Search Central Blog — Search Generative AI performance reports ↗ ; details via PPC Land ↗	JUNE 2026
Search Console Help — Generative AI performance report ↗ , worldwide since Aug. 31, 2026	CHECKED OCT. 20, 2026