

THE 2026 FIELD GUIDE FOR OWNERS AND MARKETING TEAMS

The AI Content Optimisation Playbook

Create content that machines cite and humans actually trust.

Volume is no longer a strategy. With 74% of new web pages now carrying AI-generated content, the winners are the businesses that publish what only they can prove. This playbook shows you exactly how, in five tools, backed by peer-reviewed research.

+40%

visibility in AI answers from adding statistics and sources (Princeton GEO study, KDD 2024)

4.4x

conversion rate of AI-referred visitors vs traditional organic search (Semrush, 2025)

74%

of newly published web pages already contain AI-generated content (Ahrefs, 900,000 pages)

FIVE TOOLS INSIDE

- 01 The Content Audit Framework
- 02 The Information-Gain Checklist
- 03 The Answer-First Writing Template
- 04 The Human-AI Workflow Map
- 05 The AI Share of Voice Measurement Routine

Prepared by Andrea • seoforestaurants.com

Why This Playbook Exists

Content has never been easier to produce, and that is precisely the problem. Ahrefs analysed 900,000 newly published web pages in April 2025 and found that 74.2% contained AI-generated content in some form. Only 2.5% were purely machine-written, but 71.7% were a human-AI blend, which means the blended, semi-automated article is no longer an experiment: it is the default state of the internet.

The consequence is a paradox every business leader must now navigate. Because the cost of producing ordinary content has collapsed, the value of ordinary content has collapsed with it. When any competitor can generate a plausible-looking month of blog posts before lunch, plausible-looking blog posts earn nothing. What has risen sharply in value is optimised content: material deliberately engineered to be trusted by human readers, understood by search engines, and cited by the AI assistants, ChatGPT, Gemini, Perplexity, Claude, and Google's AI Overviews, that increasingly stand between your business and your customers.

The stakes are measurable on both sides of the equation:

- The penalty for generic content is real. Google's March 2024 core update cut visibility of low-quality, unoriginal content by roughly 45%, and sites built on unreviewed, mass-produced pages lost the majority of their organic traffic within weeks. Consumer research points the same way: a Fractl / Search Engine Land study found the share of consumers who say heavy AI use would reduce their trust in a brand nearly doubled in a single year, from 20% in 2025 to 39% in 2026, and a Klaviyo/Datalily survey of 8,000 consumers found visible AI content is four times more likely to cost a brand trust than to build it (31% vs 7%).
- The reward for optimised content is equally real. The peer-reviewed GEO study from Princeton, Georgia Tech, IIT Delhi and the Allen Institute (Aggarwal et al., KDD 2024), which tested content changes across ~10,000 queries, found that adding concrete statistics, credible citations and expert quotations lifted visibility in AI-generated answers by 30–40%, the largest gains of any technique tested. And the traffic that flows from AI answers is worth more: Semrush's 2025 research found AI-referred visitors convert at 4.4x the rate of traditional organic visitors, while Ahrefs found AI referrals made up just 0.5% of its own traffic yet drove 12.1% of all signups.

THE CORE THESIS

The contest has moved from volume to credibility. You no longer win by publishing the most; you win by being the source that both a discerning human reader and a synthesising AI engine would choose over thousands of near-identical alternatives. This playbook gives you five practical tools to make that happen, in the order you should deploy them.

How to Use This Playbook

Each part is a standalone tool, but they are sequenced deliberately. Part 1 clears out the material that is quietly damaging you. Part 2 identifies the original insight only your business can publish. Part 3 shows you how to structure that insight so machines can extract and cite it. Part 4 defines where AI belongs,

and does not belong, in your production line. Part 5 gives you a simple monthly routine to measure whether it is working. A 90-day implementation roadmap and a full source list close the document.

The Landscape in Numbers

Before the tools, the evidence. These are the figures worth knowing when you make the case for this work internally, each one traceable to a named study.

Finding	Figure	Source
New web pages containing AI-generated content (Apr 2025)	74.2%	Ahrefs, study of 900,000 pages
Pages that are a human-AI blend	71.7%	Ahrefs, 2025
Visibility lift in AI answers from adding statistics, citations and quotations	30-40%	Aggarwal et al., GEO, KDD 2024 (Princeton / Georgia Tech / IIT Delhi / AI2)
Visibility lift from clarity and fluency improvements alone	15-30%	Aggarwal et al., GEO, KDD 2024
Reduction in low-quality, unoriginal content visibility after Google's core update	~45%	Google, March 2024 core update
Conversion premium of AI-referred visitors vs traditional organic	4.4x	Semrush, June 2025 (500+ topics)
ChatGPT referral conversion vs Google organic	15.9% vs 1.76%	Seer Interactive, multi-vertical study
Share of Ahrefs' signups driven by 0.5% of traffic (AI referrals)	12.1%	Ahrefs internal analysis, 2025
Position-1 organic CTR when an AI Overview is present	1.41% → 0.64%	Ahrefs, 300,000-keyword study, 2025
Consumers who say heavy AI use reduces brand trust (2025 → 2026)	20% → 39%	Fractl / Search Engine Land, 2026
Trust-less vs trust-more when consumers notice AI content	31% vs 7%	Klaviyo / Datalily, 8,000 consumers, Dec 2025

WHAT THE NUMBERS MEAN TOGETHER

The flood of automated content is not a reason to stop publishing. It is the reason distinctive publishing now wins so decisively. Algorithms are demoting the generic at scale, consumers are punishing it with their trust, and the small stream of visitors that AI engines do send arrives pre-qualified and converts at several times the rate of a blue-link click. A single citation in the right answer can quietly outperform a page-one ranking.

Part 1, The Content Audit Framework

Purpose: identify the generic material currently diluting your brand, decide what to fix, merge or remove, and establish the quality bar every future piece must clear. Most businesses discover in this exercise that 30–50% of their published library is doing nothing, or actively spending down trust.

Step 1: Build the Inventory

Export every indexable URL, blog posts, service pages, location pages, FAQs, resource downloads. For each, capture: publish date, last updated date, monthly organic sessions (last 12 months), conversions or assisted conversions, ranking keywords, and named author (or lack of one). A spreadsheet is enough; you do not need enterprise software to do this well.

Step 2: Score Every Page on Five Dimensions

Score each page 0–2 on the five dimensions below (0 = fails, 1 = partial, 2 = passes). Maximum score: 10.

Dimension	The question to ask	What a 2 looks like
Distinctiveness	Could this paragraph sit unchanged on a competitor's website?	The piece contains a point of view, example or data point only your business could publish.
Experience (E-E-A-T)	Does it demonstrate first-hand experience?	Named, credentialed author; specifics a language model could not fabricate (photos, numbers, client scenarios).
Accuracy & currency	Is every claim still true and sourced?	Claims carry verifiable sources; dates, prices and processes reflect today, not 2022.
Extractability	Could an AI engine lift a passage and cite it cleanly?	Sections open with direct answers; headings mirror real questions; one idea per passage.
Performance	Is it earning traffic, links or conversions?	Meaningful sessions, citations or conversions in the last 12 months.

Step 3: Sort Into Four Actions

Score	Verdict	Action
8–10	Asset	Keep. Refresh dates, strengthen sources, add structured data. These are your citation candidates, promote and interlink them.
5–7	Fixable	Rewrite the opening to answer-first (Part 3), inject original insight (Part 2), attach a named author, and update facts. Re-score in 60 days.
3–4	Diluting	Merge into a stronger page with a redirect, or rewrite from scratch only if the topic is commercially important.
0–2	Liability	Remove and redirect. Thin, undifferentiated pages drag down how

Score	Verdict	Action
		algorithms judge the whole domain, Google's scaled-content policies are method-agnostic and punish mass-produced pages regardless of who or what wrote them.

Step 4: Set the Bar Going Forward

Adopt a single publishing rule from this day on: no piece ships below a score of 8. This inverts the planning question. Rather than asking how many articles you can publish this quarter, ask how many pieces can genuinely demonstrate experience and add information, and let that answer set your cadence. In a market where 74% of new pages contain machine-generated content, the bar for standing out is, counter-intuitively, easier to clear than it has ever been: most of your competitors are publishing material that scores a 3.

QUICK DIAGNOSTIC

Open your three most recent blog posts. Delete your logo and brand name. If a reader could not tell which business published them, they fail the distinctiveness test, and so, increasingly, do the machines deciding whom to cite.

Part 2, The Information-Gain Checklist

Purpose: find the original insight only your business can publish. Generative engines are synthesis machines; they already contain the consensus knowledge of the public web, so content that restates it offers them nothing to cite. What they preferentially quote is information gain, what your content adds beyond what is commonly known. The GEO study found this is also where the biggest measurable lifts live: concrete statistics were the single strongest technique tested.

The Encouraging Truth About the Bar

The bar is lower than it looks. Analyses of top-ranking pages consistently find remarkably few unique data points per query, and for most searches there are valuable questions no existing page answers at all. You do not need a research department. You need to systematically publish what you already know and nobody else does.

The Checklist: Seven Sources of Original Insight

Run every planned piece through this list before drafting. If you cannot tick at least two boxes, the piece is not ready to write, it is ready to research.

✓	Source of information gain	How to capture it
☐	Your own numbers	Mine your operational data: average results across clients, before/after metrics, cost benchmarks, seasonal patterns. 'Across 40 client sites we measured X' is uncopiable.
☐	A micro-survey	Ask 50–200 customers or industry peers three questions. Even a small, honestly-reported sample produces citable statistics no one else holds.
☐	A rigorous original comparison	Test the tools, suppliers or methods your buyers compare, and publish the methodology alongside the results. Machines cite comparisons with visible method.
☐	First-hand experience	Document what actually happened: the project that went wrong, the fix that worked, the number that surprised you. Experience is the 'E' Google added to E-A-T because models cannot fabricate it convincingly.
☐	A named expert's judgement	Quote your practitioner by name and credential, taking a clear position. Quotation addition was one of the top-three techniques in the GEO study.
☐	A contrarian, defended position	Where does your experience contradict the consensus advice in your niche? A well-argued disagreement is inherently high information gain.
☐	The unanswered question	Read the top ten results for your target query and list what none of them answers. Answer that.

Packaging the Insight for Citation

- State findings as clean, self-contained sentences with numbers: 'X% of Y did Z (n = 120, March 2026).' Machines lift sentences, not vibes.
- Publish your methodology, sample size and date. Credibility markers are extraction signals.
- Give the finding its own heading phrased as the question it answers.
- Cite your own external sources properly, the GEO research found that citing reliable sources within your content increases your own likelihood of being cited.

THE COMMERCIAL LOGIC

A business that publishes its own numbers becomes the only place that information exists. AI engines are then mathematically incentivised to cite you: they cannot attribute your statistic to anyone else.

Part 3, The Answer-First Writing Template

Purpose: structure content so both readers and machines can extract the value instantly. Generative engines favour content that leads with a direct, self-contained answer and supports it afterwards, rather than burying its conclusion beneath paragraphs of preamble. Writing to be quoted, it turns out, is indistinguishable from writing clearly, the GEO study found fluency and readability improvements alone lifted visibility 15–30%.

The Template

Block	What goes there	Rules
1. Question-mirroring H1/H2	The heading phrases the exact question a buyer would ask an assistant.	Use natural language ('How much does X cost in 2026?'), not marketing labels ('Our Pricing Philosophy').
2. The direct answer	A 2–3 sentence, self-contained answer that would make sense lifted out of context.	First sentence answers the question outright, with a number or specific where possible. No throat-clearing, no 'in today's fast-paced world'.
3. The evidence	The statistics, sources, quotes and first-hand experience that earn the answer.	One idea per paragraph. Attribute every claim: named study, named person, named date.
4. The nuance	Conditions, exceptions, 'it depends' factors, handled honestly.	This is where human judgement shows and trust is built; do not hide the complexity.
5. The next step	What the reader should do, and where your business fits.	One clear action. Soft-sell; the credibility above does the selling.

A Worked Example (Opening Block)

BEFORE (TYPICAL)

In today's competitive digital landscape, more and more restaurant owners are asking themselves about online ordering commissions. There are many factors to consider, and the answer isn't always straightforward...

AFTER (ANSWER-FIRST)

Third-party delivery platforms typically charge restaurants 15–30% commission per order, and across the 25 independent restaurants we work with, moving even 20% of orders to direct channels recovered an average of £1,400 per month. Here is how those numbers break down, and when the platforms are still worth it.

Structural Signals That Support Extraction

- Attach real authorship to everything substantive: a named author with a short credential line and an author page. Anonymous content is a trust dead-end for humans and machines alike.
- Add structured data (Article, FAQPage, Organization, Author/Person schema) so the algorithm knows unambiguously what the content is and who stands behind it.
- Keep cornerstone pages genuinely current, freshness has become a visibility signal in its own right. Update the facts, then update the visible 'last updated' date; never the reverse.
- Use FAQ blocks for the secondary questions your main answer raises; each becomes an independently citable passage.

Part 4, The Human-AI Workflow Map

Purpose: define exactly where the machine should accelerate your content operation and where your experts must author. The rule the best operations codify is simple: the machine drafts, the expert transforms. Content that is merely generated is a commodity any competitor can replicate this afternoon; content that is generated, then genuinely authored, is a moat.

The Map

Production stage	Who leads	What that means in practice
Research & gap analysis	AI accelerates	Cluster topics, analyse what ranks, surface unanswered questions, summarise competitor coverage at a scale no human team could match.
Angle & point of view	HUMAN authors	The stance, the argument, the 'only we would say this', decided by your expert before a word is drafted.
Original insight & data	HUMAN authors	Your numbers, your experience, your quotes (Part 2). AI cannot supply what only you know, and must never be allowed to invent it.
Structured first draft	AI accelerates	Draft to the answer-first template (Part 3), to your brief, in your documented voice guidelines.
Expert transformation	HUMAN authors	A named human verifies every fact, injects first-hand insight the model could not have known, strips generic phrasing, and makes it sound like your brand and no one else.
Fact-check & sources	HUMAN authors	Every statistic traced to its origin; every claim attributable. This step is non-negotiable, hallucinated citations destroy exactly the credibility this whole system builds.
Consistency & QA	AI accelerates	Style checks, terminology consistency, internal-link suggestions, schema generation, accessibility passes.
Channel adaptation	AI accelerates	Adapt the finished, expert-approved piece into newsletter, social and summary formats, derivatives of authored work, not substitutes for it.
Final sign-off	HUMAN authors	The named author approves publication and puts their name on it. Accountability is a ranking asset.

Why the Division Matters Commercially

Search engines do not penalise content for being AI-assisted, Google has been explicit that its systems evaluate quality, helpfulness and trustworthiness rather than production method. What the algorithms punish, severely, is scaled content abuse: pages mass-produced primarily to manipulate rankings rather

than help a reader. The workflow above keeps you firmly on the right side of that line while capturing the genuine speed advantages of the tools. It also protects you from the consumer side of the equation: with 39% of consumers now saying heavy AI use reduces their trust in a brand, the visible layer of your content, the voice, the judgement, the experience, must be unmistakably human.

THE ONE-SENTENCE POLICY

Adopt this as written policy for your team: 'AI may draft anything; AI may publish nothing. Every piece passes through a named expert who transforms it before it ships.'

Part 5, The AI Share of Voice Routine

Purpose: measure, month by month, whether the machines advising your customers are recommending you. As answers increasingly arrive fully formed inside AI interfaces, a growing share of your content's value is delivered without a click, your brand is read, referenced and recommended inside answers your analytics never see. Ahrefs' 300,000-keyword study found position-1 click-through rates fall by more than half when an AI Overview appears. The old dashboard no longer tells the whole truth, so you add two layers to it.

Layer 1: AI Share of Voice (60–90 minutes per month)

1. Build your prompt set once: list 15–25 questions your real buyers ask, in their words: discovery questions ('best X near/for Y'), comparison questions ('X vs Y'), and problem questions ('how do I fix Z'). Keep the set stable so months are comparable.
2. On the same week each month, ask each question in ChatGPT, Gemini, Perplexity, and Google (checking the AI Overview). Use fresh sessions where possible so personalisation does not flatter you.
3. Record four things per answer in a simple spreadsheet: whether you were mentioned (yes/no); whether you were cited or linked; your position and framing (first recommendation, one of several, caveated); and who was mentioned instead.
4. Compute your monthly score: mentions ÷ (prompts × engines), tracked as a trend line. The trend matters far more than any single month's number.
5. Close the loop: for every prompt where a competitor is cited and you are not, open the cited page and identify what it has that yours lacks, usually a direct answer, a statistic, or a named source. Feed that straight into next month's content plan (Parts 2 and 3).

Layer 2: Quality of the Clicks You Do Get (30 minutes per month)

- In your analytics, segment referral traffic from chatgpt.com, perplexity.ai, gemini.google.com, copilot.microsoft.com and related domains. Many analytics setups bucket these under 'direct' or 'other' by default. Fix that first: studies suggest only around 14% of marketers currently track AI search performance at all.
- Compare conversion rate and engagement for this segment against organic search. Expect it to be small in volume and disproportionately valuable: the cross-industry benchmark is a 4.4x conversion premium (Semrush), with studies ranging from 1.3x in e-commerce to 23x in B2B.
- Report both layers together to leadership: 'mentioned in X% of buyer questions, up from Y%; AI-referred visitors converting at Zx organic.' That is the honest picture of modern visibility.

Monthly scorecard field	Example entry
Prompt set version / date	v1, 24 prompts, w/c 3 Aug 2026
Mention rate (all engines)	31% (was 24% last month)

Monthly scorecard field	Example entry
Citation rate (linked)	17%
First-recommendation rate	8%
Top competitor cited instead	[Competitor], cited in 9 answers, all pricing questions
Action generated	Publish original pricing benchmark with methodology (Part 2, box 1)
AI-referral conversion vs organic	6.1% vs 1.9% (3.2x)

Putting It Together: A 90-Day Roadmap

Phase	Weeks	Focus
Audit & clear	1–3	Run the Content Audit Framework (Part 1). Remove or redirect liabilities, shortlist your ten highest-potential fixable pages, and adopt the 8/10 publishing bar in writing.
First measurement	2	Build your prompt set and take your first AI Share of Voice baseline (Part 5). You cannot show progress without a starting number.
Insight sprint	3–6	Run the Information-Gain Checklist (Part 2). Launch one micro-survey or internal data pull; capture three expert positions on record.
Rebuild & publish	5–10	Rewrite the ten shortlisted pages to the Answer-First Template (Part 3), each carrying at least one piece of original insight and a named author. Install the Human–AI Workflow (Part 4) as standard operating procedure.
Measure & iterate	11–13	Second and third monthly AI Share of Voice checks. Feed every competitor citation gap into the next content cycle. Report the trend.

WHAT GOOD LOOKS LIKE AT DAY 90

A smaller, sharper content library where every live page clears the bar; three to five pieces carrying data that exists nowhere else on the internet; a documented workflow your whole team follows; and two months of trend data showing whether the machines advising your market have started recommending you. That is a defensible position no volume-based competitor can replicate quickly.

Sources & Further Reading

Every statistic in this playbook is drawn from the named, dated studies below.

- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., Deshpande, A., 'GEO: Generative Engine Optimization', KDD 2024 (Princeton University, Georgia Tech, IIT Delhi, Allen Institute for AI). ~10,000-query benchmark; 30–40% visibility gains from statistics, citations and quotations; 15–30% from fluency.
- Ahrefs (2025), 'What Percentage of New Content Is AI-Generated?' Analysis of 900,000 newly published pages, April 2025: 74.2% contained AI content; 2.5% pure AI; 71.7% blended.
- Ahrefs (2025), 300,000-keyword AI Overviews CTR study: position-1 organic CTR 1.41% → 0.64% where an AI Overview appears; internal analysis showing 0.5% AI-referral traffic driving 12.1% of signups.
- Semrush (June 2025), AI Search traffic study across 500+ topics: AI-referred visitors convert at 4.4x traditional organic.
- Seer Interactive (2025–26), Multi-vertical referral study: ChatGPT referrals converting at 15.9% vs 1.76% for Google organic.
- Google (March 2024), Core update and scaled content abuse policy: ~45% reduction in low-quality, unoriginal content in results; policy explicitly method-agnostic.
- Fractl / Search Engine Land (2026), Consumer survey (1,008 US consumers): share saying heavy AI use would reduce brand trust rose from 20% (2025) to 39% (2026).
- Klaviyo / Datalily (Dec 2025), Survey of 8,000 consumers across 8 countries: visible AI content makes 31% trust a brand less vs 7% more.
- Pew Research Center (2025), Browsing study of 900 US adults: clicks on traditional results fall from 15% to 8% on pages showing an AI summary.

Figures reflect the studies as published; ranges and effects will vary by industry and competitive context. Treat them as directional benchmarks, not guarantees.

Ready to Publish Content That People and Algorithms Both Believe?

If you have worked through this playbook, you already understand the ground has shifted permanently. The flood of automated content will only rise, and the engines deciding who gets recommended will only grow more selective. For a business willing to do this properly, that is not a threat. It is the clearest competitive opening in years, because the harder distinctive content becomes to fake, the more decisively it wins.

LET'S TALK IT THROUGH

If you would like to talk more about any part of this playbook, or you are interested in what it would look like applied to your business specifically, we would love to invite you to book a call with Andrea. The call is free and focused, and you will leave with a candid read on how your current content is being received by your customers and by the machines advising them.



[Book your free strategy call with Andrea](https://seoforestaurants.com/booking)

seoforestaurants.com/booking

Until next time: keep creating with intent and publishing with confidence. The businesses that give the machines something worth quoting, and their audiences something worth believing, are the ones quietly becoming the names their market recommends first.