

ASQ · INFORMATION QUALITY SESSION · 2026

When AI Search Becomes the Company's **Front Door**

From SEO to AEO – and the new risks to information quality when a machine answers for you.

A quality-professional's field guide to AI Search Optimization

Karla Lawson · VP of Digital Marketing

YOUR SPEAKER



Karla Lawson

VP of Digital Marketing

- 15+ years leading digital marketing across SaaS, financial services, and homebuilding.
- Depth in SEO / AEO, website replatforming, MarTech, and analytics.
- Brings a quality lens to AI search: defect rates, monitoring, and governance.

Email karla.lawson@me.com

LinkedIn linkedin.com/in/lawsonkarla

What you'll leave with

Four ideas, in language quality professionals already use.

01

The front door moved

Why AI answers are replacing your website as the first thing people learn about you.

02

It's a defect-rate problem

Miscitation, fabrication and non-reproducibility - measured, not hypothetical.

03

You can't inspect the output

So you shift left: govern the inputs the machine reads instead of the answer it writes.

04

It's a QMS problem

Monitoring, ownership and governance - not a one-time marketing campaign.

THE SHIFT

**Your reputation used to live on
surfaces you controlled.
Now it lives inside an answer **you
never wrote.****

The buyer, the auditor, the candidate - they increasingly get a synthesized answer before they ever reach a page you QC'd. The source list is assembled on the fly, mostly from properties you don't own.

YESTERDAY

- Search returns your page.
- You control the words.
- You inspect before it ships.

TODAY

- AI writes the answer.
- Many sources, mostly not yours.
- No inspection step at all.

FOUNDATIONS

SEO, AEO, GEO: three letters, one shift

Search stopped just pointing to answers and started writing them. Same goal - get found - but the target moved.

SEO

Search Engine Optimization

Optimize to rank a link on the results page.
The reader clicks through to you.

Target: a position

AEO

Answer Engine Optimization

Optimize to be the answer - snippets, voice, AI answer boxes. The engine states it; the click is optional.

Target: the answer

GEO

Generative Engine Optimization

Optimize to be used and cited accurately inside an AI-generated answer synthesized from many sources.

Target: the citation

They're layers, not replacements. SEO still drives the most traffic today - but AEO and GEO decide whether the machine represents you at all. AEO and GEO are at times used interchangeably

FOUNDATIONS

Five realities of AI search

Why this behaves nothing like the search you optimized for a decade.

● Non-deterministic

No two runs return the same answer - variability is built in.

● You don't own the inputs

It draws on reviews, forums and press - channels outside your control.

● Cross-functional

It needs SEO, content, PR and data pulling the same direction.

● Gaming won't last

Shortcuts get suppressed over time, the way SEO spam did.

● Competitors are moving

Rivals are setting their AI baselines right now.

● The quality lens

A process with many uncontrolled inputs and high common-cause variation - manage and monitor it, don't set-and-forget.

FOUNDATIONS

How AI forms a view of your brand

AI blends signals from everywhere into one answer. Three signal types, three jobs - and only one you fully control.

EARNED

What others say

Reviews, articles, social posts, forums.

Builds trust. Third-party credibility carries the most weight with AI.

OWNED

What you publish

Your site, docs, structured data.

Provides clarity - the accurate facts. The only bucket you fully control.

PAID

What you place

Ads and sponsored placements.

Reinforces the message across the broader web.

Why it matters: *the model synthesizes all three, but you control OWNED outright - which is where a quality team has the most leverage*

FOUNDATIONS

Where to focus: the sources AI cites for B2B

Most AI citations point to third-party sites, not your own. For B2B software the pool narrows to ~15 recurring sources - so concentrate effort where the machine actually looks.

● REVIEW & COMPARISON

G2 – most-trusted review site
Capterra – software comparisons
Software Advice – vendor evaluations
TrustRadius – verified reviews

● COMMUNITY & UGC

Reddit – often #1 for Perplexity
YouTube – demos; transcripts ingested
LinkedIn – company profiles
Quora – nuanced B2B answers

● REFERENCE & NEWS

Wikipedia – neutral company facts
Gartner / Forbes – expert analysis
PCMag / TechRadar – top-10 listicles
TechCrunch – industry news

Quality lens: you don't own these sources and can't inspect them line-by-line - so earning accurate, consistent presence on them is the work, not a nice-to-have.

FOUNDATIONS

The engines don't all think alike

Where each one looks for answers differs - so who gets cited, and where you focus, shifts by platform.

● ChatGPT

The Authority All-Rounder

Favors established knowledge bases over social chatter - Wikipedia, authoritative tech and news media - with a strong pull from G2, Capterra and Gartner. Limited Reddit.

● Perplexity

The Community Specialist

Leans hard on Reddit and niche forums (upvotes as a relevance proxy), YouTube for walkthroughs and Gartner for framing, plus LinkedIn, G2 and PCMag.

● Claude

The Structured Intellectual

Deep-dive blogs, white papers, official docs and vetted SaaS marketplaces. Minimal forum reliance; rewards rigor, structure and logical consistency.

● Gemini

The Ecosystem Player

Google Search results, high-authority domains and Google-owned properties, plus Reddit for perspective, YouTube for how-to, and LinkedIn / Quora for expert input.

02

ACT II

It's an information-quality problem

Miscitation, confident fabrication, and a measurement system that doesn't reproduce.

The defect rate is measured, not hypothetical

>60%

of AI-search answers cited news incorrectly across 8 engines (1,600 queries).

Tow Center / CIR, 2025

94%

error rate for the worst engine tested (Grok-3) - premium tools were confidently wrong.

Tow Center / CIR, 2025

~1 in 3

brand mentions get a fact wrong when AI does mention you.

Webflow audit / CMSWire, 2026

16% / 6%

the median brand appears in only 16% of relevant answers; cited just 6% of the time.

Webflow audit / CMSWire, 2026

THE FAILURE MODE

The risk isn't silence. It's a confident wrong answer.

When your content is thin or ambiguous, the model doesn't leave a blank - it fills the gap with a plausible guess and presents it with full confidence.

These systems are bad at declining to answer. They rarely say “I don't know.”

In quality terms

- A gauge that is merely noisy is a nuisance.
- **A gauge that never flags out-of-spec and confidently reports the wrong value is a hazard.**
- That's the AI-answer failure mode: no abstention, no error bar, delivered at the moment a decision gets made.

Ask the same question twice, get two different answers

2,961 identical prompts run across ChatGPT, Claude and Google AI (600 volunteers):

< 1 in 100

chance the same brand list appears twice

~ 1 in 1,000

chance the list repeats in the same order

This is a measurement-system problem

A “rank #1 in ChatGPT” number is one draw from a distribution, not a stable value. Reported as a trend line to leadership, it fails reproducibility the way a gauge fails a Gauge R&R study. Before you act on any AI-visibility metric, ask whether the measurement itself reproduces.

THE SIGNAL INSIDE THE NOISE

Rank is noise. Appearance frequency is a real signal.

Individual answers vary every time - but run the prompt 50-100 times and a stable distribution emerges. Some brands appear in most responses; most appear rarely. That frequency is measurable and trackable.



Track the distribution, not the draw

Appearance rate over many runs is stable where any single rank is not.



Brand Visibility Score

Answers mentioning you $\div$ total answers for your space $\times$ 100. A capability metric, not a rank.



What drove inclusion

Top-appearing brands weren't better optimized - they had deeper earned media in trusted sources.

03

ACT III

Govern the inputs, monitor the process

You can't inspect the AI's output - so shift quality left and control what it reads.

Four inputs you can actually control

Classic shift-left: build quality into the source material the model ingests.

1

Source consistency

Say the same facts the same way everywhere - owned pages, directories, reviews, press. Contradictions are what the model turns into fabrication.

2

Structured data / schema

Machine-readable markup that tells engines exactly what each page is. Source-side quality control for how you're parsed.

3

Authoritative human content

Original expertise and credible sourcing (E-E-A-T). Thin or purely-AI content gets low quality signals - and low trust.

4

One canonical entity

A single, clearly-defined “organization” that every page references, so the machine resolves you to one clean identity.

CASE STUDY

Making an enterprise B2B brand legible to AI

A real GEO-readiness engagement - that shows the problem and the fix on one project.

● WHO

A global enterprise software company.
In AI answers, a buyer meets a synthesized description of the brand before ever reaching its site.

● THE STAKES

For B2B software, AI recommendations are assembled mostly from third-party sources. The company's own read: competitors were setting baselines faster.

● THE MOVE

The company already tracks AI visibility through third-party tools and brought in a digital-experience partner to audit the signals it owns.

CASE STUDY · THE PROBLEM

The defect was in the one input they control

You can't fix what third parties say - but your own structured data is the signal you own. The company's was mixed, and getting worse.

Product snippets breaking

700+ invalid vs 50 valid; 4 critical issues; invalid trend rising

AI impact: *The volume of broken product signals was growing, not shrinking.*

Homepage overextended

The corporate entity was declared twice, alongside invalid Product and Video markup.

AI impact: *The cleanest signal for “who the company is” was the muddiest.*

Inconsistent software model

Some pages used SoftwareApplication, some Product, some just WebPage; a flagship product page carried an invalid Offer.

AI impact: *The machine couldn't tell what a product page even was.*

Resources untyped

Reports, case studies and videos exposed only generic Organization schema.

AI impact: *Assets weren't legible as report, case study or video.*

Make the entity model unambiguous - then earn authority

Fix the plumbing (schema)

- One canonical Organization every page references - no duplicate entities.
- Named software → SoftwareApplication, connected back to that organization.
- Resource schema driven by real asset type - DigitalDocument, Report, Case Study, VideoObject - not the URL or a generic template.

Earn the trust (content)

- Own the topic, not just keywords - hub-and-spoke coverage of core categories.
- Build brand mentions in the third-party sources AI trusts - now weighted above backlinks.
- Human E-E-A-T content with real expertise, not thin AI output.
- Measure representation with the existing stack (Evertune, SEMRush, Console, Webmaster Tools, Ahrefs, GA4, etc).

The point: *clean signals make you parseable; authority makes you citable. The company needed both - the schema fix without the content play just makes you accurately ignored.*

A phased plan that ends in governance, not a cleanup

NOW

Fix visible defects

Homepage cleanup, software pages →
SoftwareApplication, blog / news date
QA, VideoObject where metadata exists.



NEXT

Extend the templates

Resources by asset type, Customer
Success → service catalog, industry pages
connect to #organization.



FUTURE

Govern recurrence

One canonical Organization, an @id
reference rule, a CMS schema decision
tree.

How they'll know it worked

This is a remediation plan, not a results report. Success shows up later as fewer invalid items and cleaner AI mentions - tracked in their monitoring stack. The durable win is the last phase: governance keeps the defects from growing back.

Does optimizing actually move the needle?

What the research found

up to +40% visibility in generated answers from adding statistics, citations and quotes from credible authorities.

Keyword density - the old SEO lever - barely mattered.

Aggarwal et al., KDD 2024 (GEO-bench, ~10,000 queries)

Read it like a quality engineer

That +40% is a **maximum relative gain in a simulator** (five sources pre-placed, one model), not a promise of ranking first in ChatGPT.

The discipline is young, its metrics are contested, and long-run efficacy is unproven. Vendors routinely overstate this figure.

Pilot and measure on your own content before you believe anyone's number.

You can't do final inspection - so monitor the process

There's no incoming-inspection gate on a generated answer. Treat AI representation like a process you chart over time, not a page you sign off once.

Appearance rate

How often you show up across many runs of a prompt.

Accuracy rate

Share of mentions that state your facts correctly.

Sentiment & framing

How you're characterized, not just whether you're named.

Source mix

Which third-party properties are feeding the answer.

Caution on tools. Monitoring vendors exist (Evertune, Ahrefs, Semrush, Brightedge, Similarweb, Google Console & Bing Webmaster and others). Their numbers are useful for trend, but many rest on the same shaky “rank” assumption without full datasets. Validate methodology - especially reproducibility - before you report a vendor metric upward as fact.

Monday Morning Quality Check

1

Baseline the defect rate

Ask ChatGPT, Claude and Google what your org does and sells. Log what's wrong. That's your incoming quality report.

2

Fix source-side defects

Correct contradictions, clean up schema, resolve to one canonical entity. Remove what the model turns into fabrication.

3

Invest in authority

Earned media and genuine human expertise. Brand mentions in trusted sources now outweigh backlinks for inclusion.

4

Stand up monitoring

Track appearance and accuracy over time, not a rank. Chart it like any process metric.

5

Assign an owner

Name who is accountable. This spans SEO, content, PR and data - so it needs governance, not goodwill.

THE TAKEAWAY

“Clarity is the best defense against distortion.”

- Forbes Technology Council, on AI search

- The front door moved from pages you inspect to answers you don't.
- It's a defect-rate and reproducibility problem - quality's home turf.
- You can't inspect the output, so govern the inputs and monitor the process.

Thank you · Questions?

Sources & notes

Tow Center for Digital Journalism / CJR (2025)

"AI Search Has a Citation Problem." 8 engines, 1,600 queries; >60% incorrect; Grok-3 94%.

Webflow AEO audit via CMSWire (2026)

2,000+ US sites across ChatGPT/Gemini/Perplexity/Claude; ~1/3 of mentions inaccurate; 16% appearance, 6% citation. Vendor research - treat as directional.

Forbes Technology Council (Apr 2026)

"Beyond the Citation: Managing Brand Risk in the Age of AI Search." Fabrication as the real risk.

SparkToro / Gumshoe.ai - Fishkin & O'Donnell (Jan 2026)

2,961 prompts; <1-in-100 same list, ~1-in-1,000 same order; appearance frequency is the stable signal.

Aggarwal et al., KDD 2024 (arXiv:2311.09735)

"GEO: Generative Engine Optimization." GEO-bench ~10k queries; authority signals +≈40% (simulator max).

Critical Survey of GEO 2023-2026 (arXiv, 2026)

Cautions that the +40% figure is often overstated; measurement contested.

Anonymized GEO schema audit (2026)

Real-world schema case study (client anonymized) - foundation strong, page-level defects, phased remediation.

Enterprise B2B LLM-search review (2026)

Top B2B citation sources; earned/owned/paid signal mix; per Higoodie domain analysis.

Companion pre-reading also included Fishkin's SparkToro video ("Why AI Rankings Don't Exist") and CMSWire's coverage.