



What actually changes in SEO when the goal expands from ranking in traditional Google search results to being retrieved, cited, mentioned, or recommended in AI-generated answers such as Google AI Overviews or AI Mode, ChatGPT, Gemini and Perplexity?

Distinguish what remains foundational SEO from what genuinely changes for AI Search. Use current evidence and cite your sources.

When the goal expands from ranking in traditional Google results to being retrieved, cited, or recommended inside AI-generated answers (Google AI Overviews/AI Mode, ChatGPT, Gemini, Perplexity), the foundation of SEO does not disappear—but the optimization target, success metrics, and some tactical priorities do shift. ^{[1] [2] [3]}

What stays foundational (still SEO)

Google and other major players explicitly state that optimizing for generative AI surfaces is still SEO, because these systems lean on the same core index and quality signals. In practice, the following remain essential: ^{[2] [1]}

- **Crawlability and indexation:** Pages must be discoverable, allowed in robots.txt, and indexable. AI systems still need to fetch and understand your content. ^{[4] [5] [6]}
- **Technical health:** Core Web Vitals, mobile-friendliness, HTTPS, clean HTML, and proper site architecture still matter. ^{[7] [8]}
- **E-E-A-T and authority:** Expertise, experience, authoritativeness, and trustworthiness remain central to whether your content is considered reliable enough to be cited. ^{[3] [9] [10]}
- **Entity clarity and structured data:** Clear entities (people, organizations, products) and schema (Article, FAQPage, Organization, Person, BreadcrumbList, etc.) help both traditional ranking and AI extraction. ^{[11] [4] [7]}
- **Intent-matching content:** Pages must still satisfy user intent and cover topics comprehensively; thin or scaled content is penalized in both worlds. ^{[9] [4]}

Google's own guidance frames AEO (Answer Engine Optimization) and GEO (Generative Engine Optimization) as extensions of SEO, not replacements. ^{[1] [2]}

What genuinely changes for AI search (GEO/AEO)

The shift is from “rank a URL in a list” to “become a source the model chooses to quote or recommend.” That changes the unit of optimization and the KPIs. ^[10] ^[12] ^[13]

1) Optimization target: from position to presence

- **Traditional SEO:** Compete for positions 1–10 in a ranked list; success is measured by rankings, impressions, CTR, and organic sessions. ^[13] ^[14]
- **AI search (GEO/AEO):** Compete for **in-answer presence**—whether your brand/page is cited inside the AI summary at all. There is no “position 2” inside an overview; it’s cited or not. ^[14] ^[10]

Practically, this means designing content so an AI can **extract a concise, self-contained answer** and attribute it to you. ^[4] ^[7] ^[11]

2) Content structure: answer-first, snippet-ready

AI systems favor content that is easy to retrieve, validate, and lift into a response. Common patterns seen in 2025–2026 guidance:

- **Lead with a direct answer:** Place a 40–60 word, standalone answer immediately under each major heading. ^[15] ^[11] ^[4]
- **Self-contained passages:** Write sections that make sense out of context (roughly 130–170 words per key passage) so they can be quoted without extra context. ^[16] ^[17]
- **Explicit Q&A and lists:** Use clear questions as headings and support with bulleted/numbered evidence; apply FAQPage schema where appropriate. ^[6] ^[11] ^[4]
- **Source-backed facts:** Attribute statistics and claims to named sources with links; AI engines weigh provenance and cross-source agreement. ^[11] ^[13] ^[4]

This is different from traditional long-form “everything on one page” SEO, which often buries the answer deep in the copy. ^[13] ^[16]

3) Signals that matter more for being cited

While backlinks and domain authority still help, AI citation decisions additionally reward:

- **Factual precision and uniqueness:** Specific, dated data points and original analysis increase the chance of being selected as a source. ^[10] ^[13]
- **Entity coverage and semantic completeness:** Thoroughly covering sub-questions and related entities makes your page a better “knowledge source” for multi-part queries. ^[17] ^[10] ^[13]
- **Freshness:** For time-sensitive topics, newer or recently updated content is favored, especially by retrieval-heavy engines like Perplexity. ^[18] ^[4]
- **Provenance and trust signals:** Clear authorship, organizational identity, and consistent entity markup across the web improve trust in your content as a citation. ^[10] ^[13]

4) Bot access and platform-specific nuances

AI search is fragmented across engines with different crawlers and retrieval logic:

- **Robots.txt for AI bots:** Explicitly allow relevant agents (e.g., GPTBot, OAI-SearchBot, PerplexityBot, ClaudeBot, Google-Extended) on pages you want cited. ^[5] ^[6] ^[11]
- **Engine differences:**
 - **Google AI Overviews / AI Mode:** Heavily tied to Google's existing index; traditional SEO strength, structured data, and topical depth are crucial. ^[8] ^[6]
 - **Perplexity:** Favors fresh, specific facts and plain statements of data; heavily citation-weighted and retrieval-driven. ^[6] ^[18] ^[4]
 - **ChatGPT (Search):** Requires allowing OAI-SearchBot and structuring content for easy extraction (tables, lists, self-contained paragraphs). ^[5] ^[6]

This means you may need **platform-aware tweaks** (e.g., freshness cadence for Perplexity, explicit bot permissions for OpenAI) beyond generic SEO. ^[18] ^[5]

5) New KPIs and measurement

Because the win condition changes, so do the metrics:

- **Traditional SEO KPIs:** Rankings, organic traffic, CTR, conversions from search. ^[14] ^[13]
- **AI search / GEO KPIs:** Citation rate, “in-answer presence,” visibility score, share of voice in AI answers, brand trust lift. ^[19] ^[13] ^[14]

These are not available in Google Search Console and often require specialized tracking or manual auditing of AI answers. ^[14]

A practical way to think about it

- **Foundational SEO** = make your site crawlable, fast, authoritative, and intent-aligned so it's eligible to be considered. ^[2] ^[1]
- **AI search optimization (GEO/AEO)** = make your facts, entities, and answers easy for models to retrieve, validate, and cite, then measure success by **citation share** rather than just rank. ^[12] ^[19] ^[10]

In 2026, the most effective approach is hybrid: keep strong traditional SEO as the base, then layer answer-first structure, explicit sourcing, entity clarity, and AI-bot access to improve your odds of being quoted in AI-generated answers. ^[3] ^[7] ^[9]



Looking specifically at the changes you just identified, which are:

verified by platform documentation or strong experimental evidence, supported mainly by observational correlation or reasonable inference, and commonly promoted by the SEO/GEO industry without strong evidence? For each important claim, show the evidence and source. Do not treat correlation as causation.

Here's how the main claims about "what changes for AI search" break down by strength of evidence as of mid-2026.

1. Verified by platform documentation or strong experimental evidence

A. "Optimizing for AI Overviews / AI Mode is still SEO" (Google)

Claim: Google's AI Overviews and AI Mode use the same index and ranking systems as traditional Search; there is no separate "AI ranking algorithm," and no special AI-only requirements.

Evidence:

- Google's official guide "Optimizing your website for generative AI features on Google Search" (May 15, 2026) states that AI Overviews and AI Mode are built on top of Google's core Search systems and that "there are no additional requirements to appear in AI Overviews or AI Mode, nor other special optimizations necessary."^[31] ^[32] ^[33]
- Multiple summaries of that document confirm Google's position that AEO and GEO are "still SEO," not separate disciplines.^[34] ^[35] ^[36]

Status: Verified by platform documentation.

B. "No special llms.txt, content chunking, or AI-specific markup needed for Google"

Claim: For Google Search (including AI Overviews and AI Mode), you do not need:

- `llms.txt` files,
- artificially chunked content,
- AI-specific schema or rewritten "AI versions" of pages.

Evidence:

- Google's May 2026 AI optimization guide explicitly lists these as unnecessary: `llms.txt`, content chunking, special schema for AI, and AI-only rewrites.^[33] ^[37] ^[38] ^[39] ^[40] ^[41]
- The guide explains that Google's systems can understand multi-topic pages and extract relevant passages without pre-fragmentation.^[37] ^[38]

Status: Verified by platform documentation.

C. “AI Overviews draw from the same Search index; E-E-A-T still central”

Claim: AI Overviews pull from the same Google index as organic results and rely on the same quality signals (including E-E-A-T).

Evidence:

- Google’s guidance and multiple analyses state there is no separate “AI index”; AI Overviews use the same index and quality frameworks as traditional Search. ^[35] ^[36] ^[42] ^[33]
- The guidance emphasizes people-first, high-quality content and standard structured data as what surfaces in AI Overviews. ^[42]

Status: Verified by platform documentation (for the “same index, same quality framework” part). Exact weighting of E-E-A-T inside AI Overviews is not fully disclosed, but the principle that quality signals matter is explicit.

D. “OpenAI has distinct crawlers: GPTBot (training) vs OAI-SearchBot (ChatGPT Search citations)”

Claim:

- GPTBot is for training OpenAI’s foundation models.
- OAI-SearchBot is for indexing pages so ChatGPT can retrieve and cite them in search answers.
- You can control each independently via `robots.txt`.

Evidence:

- OpenAI’s official crawler documentation (platform.openai.com/docs/bots) describes both bots, their user agents, IP ranges, and behavior, and confirms `robots.txt` support with ~24-hour propagation. ^[43] ^[44] ^[45] ^[46] ^[47]
- Multiple technical guides correctly distinguish the two bots and show working `robots.txt` examples. ^[44] ^[46] ^[43]

Status: Verified by platform documentation.

E. “Perplexity uses live retrieval and reranks on freshness, relevance, authority, structure”

Claim: Perplexity:

- Performs fresh web retrieval for queries (no static answer store).
- Uses a multi-stage pipeline: retrieval → filtering → reranking based on factors like freshness, semantic relevance, authority, and structure.
- Has documented bots (PerplexityBot, Perplexity-User) for indexing and live fetches.

Evidence:

- Perplexity’s help docs and multiple technical explainers describe real-time retrieval with no fixed knowledge cutoff and a reranking system that weighs freshness, relevance, and authority. ^[48] ^[49] ^[50] ^[51]
- Independent analyses outline a 5-stage pipeline with gates for freshness, semantic relevance, engagement, and crawl access. ^[49] ^[50] ^[51] ^[52] ^[53]

Status: Strongly supported by platform documentation + consistent technical reporting. Exact weights are not public, but the existence of these factors is well documented.

2. Supported mainly by observational correlation or reasonable inference

These claims have data but are not proven causal by controlled experiments or official statements.

A. “Answer-first / BLUF / self-contained passages increase AI citations”

Claim: Structuring content with:

- A direct 40–60 word answer at the start of sections,
- Self-contained 120–180 word passages per sub-question,
- Clear Q&A headings and lists,

makes pages more likely to be cited in AI answers.

Evidence:

- Many GEO/AEO playbooks recommend “answer first” and “self-contained passages” based on how RAG systems retrieve and quote text. ^[54] ^[55] ^[56] ^[57] ^[58] ^[59]
- Perplexity-focused analyses note that BLUF-style, directly responsive content tends to rank higher in their reranking because of semantic relevance. ^[51] ^[52] ^[60]

But:

- There is no widely cited, peer-reviewed or large-scale controlled experiment showing that adding answer-first structure *causes* higher citation rates, holding other factors constant.
- Most evidence is observational (pages that are cited often already look like this) or inferred from how retrieval systems work.

Status: Observational correlation + plausible mechanism, not experimentally verified.

B. “Schema/structured data correlates with higher AI citation rates”

Claim: Pages with schema markup (FAQ, Article, HowTo, etc.) are more likely to be cited in AI Overviews, ChatGPT, Perplexity, etc.

Evidence for correlation:

- Multiple industry analyses report positive correlations:

- “Pages with schema markup are 36% more likely to appear in AI-generated citations” (WPRiders). ^[61]
- “89% correlation between valid schema and Perplexity citations” (AISEO.com.mx, 650+ sites). ^[61]
- Princeton/Moz-style study reporting FAQ schema pages having ~3.2× the AI Overview citation rate of pages without schema. ^[62]
- Cross-sectional studies consistently find that cited pages are more likely to carry schema than non-cited pages. ^{[63] [62] [61]}

Evidence against causation:

- Ahrefs’ 2026 controlled experiment (1,885 pages adding JSON-LD vs 4,000 matched controls) found **no meaningful uplift** in AI citations after adding schema:
 - Google AI Overviews: −4.6% (small but statistically significant decline),
 - AI Mode: +2.4% (not distinguishable from zero),
 - ChatGPT: +2.2% (not distinguishable from zero). ^{[64] [65] [66] [67] [68] [63]}
- Other analysts conclude: “Schema markup correlates with AI citations but may not cause them.” ^{[67] [68] [63]}

Status: Correlation is real; causation is not supported by the best controlled evidence. The most defensible statement is: “Pages that are already strong tend to have schema; adding schema alone doesn’t reliably increase citations.”

C. “Brand mentions correlate strongly with AI citation / visibility”

Claim: The volume of brand mentions across the web (linked or unlinked) is a strong predictor of whether a brand appears in AI Overviews and other AI answers.

Evidence:

- Ahrefs’ study of ~75,000 brands found:
 - Brand web mentions correlate with AI Overview visibility at $r \approx 0.664$, about 3× stronger than backlinks (referring domains at ~0.218). ^{[69] [70] [71]}
 - Top-quartile brands by mentions average far more AI Overview mentions than lower tiers. ^[71]
- Other analyses report similar patterns: strong association between brand visibility and AI citation. ^{[72] [69]}

But:

- These are **correlational studies**; they do not prove that increasing mentions *causes* more AI citations (though it’s plausible via entity recognition and trust signals).
- No large-scale controlled experiment isolates “brand mentions” as the only changed variable.

Status: Strong observational correlation; plausible causal story, but not experimentally verified.

D. “Being cited in AI Overviews increases clicks / brand value”

Claim: Brands cited inside AI Overviews get more organic (and sometimes paid) clicks than uncited brands on the same queries.

Evidence:

- Seer Interactive (2026) reports that brands cited in AI Overviews earn approximately 120% more organic clicks per impression than uncited brands on the same query. ^[73]
- Other analyses show cited brands getting ~35% more organic clicks and ~91% more paid clicks than non-cited competitors on the same SERP. ^[74]

But:

- These are observational comparisons; they don’t fully control for pre-existing brand strength, query intent, or position effects.
- They show association, not a clean causal effect of “being cited → more clicks.”

Status: Observational correlation with strong business implication; causation not fully isolated.

E. “Freshness and topical authority matter more for AI citations than for traditional rankings”

Claim: For AI engines (especially Perplexity and real-time retrieval systems), freshness and topical authority are relatively more important than in classic ranking.

Evidence:

- Perplexity documentation and analyses explicitly list freshness and topical authority among top reranking factors. ^{[52] [60] [48] [49] [51]}
- Industry reports on AI search note that time-sensitive queries heavily favor newer content in AI answers. ^{[75] [54]}

But:

- Google has not published a detailed breakdown showing that freshness is *more* important in AI Overviews than in regular Search; it’s already a known ranking factor.
- Most “more important” claims are inferred from observed behavior and platform descriptions, not side-by-side controlled comparisons.

Status: Reasonable inference + platform descriptions; not rigorously quantified vs traditional ranking.

3. Commonly promoted by the SEO/GEO industry without strong evidence

These are widely marketed tactics or narratives that lack solid experimental or official support, and in some cases are explicitly downplayed by platforms.

A. “You must create `llms.txt` to be cited by AI (including Google)”

Claim: Publishing an `llms.txt` file is necessary or highly beneficial for appearing in AI Overviews, ChatGPT, Perplexity, etc.

Evidence:

- Google’s May 2026 AI optimization guide explicitly says `llms.txt` is **not needed** for AI Overviews or AI Mode; Google’s systems don’t treat it specially. ^[38] ^[40] ^[41] ^[33] ^[37]
- Independent checks report that major AI crawlers (Google, OpenAI, Anthropic) do not rely on `llms.txt` for indexing or citation decisions. ^[38]

Reality:

- Some non-Google tools *might* read `llms.txt`, but for Google’s AI features it is **not a lever**.
- Many vendors still promote it as a core GEO tactic despite Google’s explicit guidance.

Status: Overstated / not supported for Google; at best a niche signal for some third-party tools.

B. “You must artificially chunk content into small pieces for AI to understand/cite you”

Claim: Breaking articles into tiny, AI-optimized chunks is required to get cited in AI Overviews or other AI answers.

Evidence:

- Google’s AI optimization guide states there is **no need** to break content into small pieces; its systems can understand multi-topic pages and extract relevant passages. ^[39] ^[33] ^[37] ^[38]
- At Search Central Live Milan, Google reiterated that content chunking for AI extraction “does not help.” ^[39]

Reality:

- Logical, well-structured sections are good for humans and may help retrieval, but **artificial micro-chunking** as a GEO requirement is not supported by Google and lacks controlled evidence.

Status: Industry hype beyond what platforms or experiments support.

C. “Adding schema will reliably increase your AI citations”

Claim: If you add FAQ, Article, or other schema, you will see a clear uplift in citations from Google AI Overviews, ChatGPT, Perplexity, etc.

Evidence:

- As noted above, cross-sectional studies show correlation, but the **best controlled experiment** (Ahrefs, 2026) found **no meaningful uplift** and even a small decline for AI Overviews after adding schema. ^{[65] [66] [68] [63] [64] [67]}
- Google’s guidance says no special schema is required for AI Overviews or AI Mode. ^{[37] [42] [39]}

Reality:

- Schema is still valuable for classic rich results and clarity, but marketing it as a guaranteed AI-citation booster is not supported by controlled data.

Status: Overpromoted relative to experimental evidence.

D. “GEO/AEO are entirely new disciplines separate from SEO”

Claim: You need a completely new playbook, separate from SEO, to win in AI Overviews and AI search.

Evidence:

- Google explicitly frames AEO and GEO as “**still SEO**”, not separate disciplines, and integrates them into Search Central’s existing guidance. ^{[32] [36] [31] [34] [35]}
- Many reputable analysts echo this: foundational SEO is the base; GEO/AEO are extensions, not replacements. ^{[76] [77] [78] [79]}

Reality:

- The *emphasis* shifts (e.g., more focus on citations, brand mentions, answer structure), but the core is still SEO.
- Framing GEO/AEO as wholly new is often a marketing narrative rather than a technical necessity.

Status: More marketing framing than technical reality, per Google’s own stance.

E. “If you’re not in AI Overviews, your SEO is dead / you must completely rebuild”

Claim: AI Overviews make traditional SEO obsolete; you must pivot entirely to GEO or lose all traffic.

Evidence:

- Data shows AI Overviews reduce CTR on some informational queries, but organic search still drives substantial traffic, and many queries still show classic results. ^{[80] [81] [82] [83]}

- Google’s guidance emphasizes continuing to invest in high-quality, people-first SEO; it does not suggest abandoning traditional practices. [\[36\]](#) [\[31\]](#) [\[32\]](#)

Reality:

- AI Overviews change the *mix* of visibility and clicks, but they do not eliminate the value of ranking well in organic results.
- “SEO is dead” narratives are recurrent and not supported by traffic data or platform guidance.

Status: Fear-based marketing, not evidence-based.

Summary map

Verified by platform docs / strong experiments:

- AI Overviews/AI Mode use the same index and quality framework as Google Search; no special AI-only requirements. [\[31\]](#) [\[32\]](#) [\[33\]](#) [\[35\]](#) [\[36\]](#) [\[42\]](#)
- `llms.txt`, artificial chunking, AI-specific schema/rewrites are **not needed** for Google’s AI features. [\[40\]](#) [\[41\]](#) [\[33\]](#) [\[37\]](#) [\[38\]](#) [\[39\]](#)
- OpenAI’s GPTBot vs OAI-SearchBot distinction and `robots.txt` control are documented. [\[45\]](#) [\[46\]](#) [\[47\]](#) [\[43\]](#) [\[44\]](#)
- Perplexity uses live retrieval with reranking on freshness, relevance, authority, structure; bots are documented. [\[50\]](#) [\[53\]](#) [\[48\]](#) [\[49\]](#) [\[51\]](#) [\[52\]](#)

Observational correlation / reasonable inference:

- Answer-first, self-contained passages correlate with higher AI citations; plausible but not experimentally proven. [\[55\]](#) [\[56\]](#) [\[57\]](#) [\[59\]](#) [\[51\]](#) [\[52\]](#) [\[54\]](#)
- Schema correlates with AI citations, but controlled tests show **no causal uplift** from adding schema. [\[66\]](#) [\[68\]](#) [\[62\]](#) [\[63\]](#) [\[64\]](#) [\[65\]](#) [\[67\]](#) [\[61\]](#)
- Brand mentions strongly correlate with AI visibility; causation plausible but not isolated. [\[70\]](#) [\[69\]](#) [\[71\]](#) [\[72\]](#)
- Being cited in AI Overviews associates with higher clicks; causation not fully isolated. [\[73\]](#) [\[74\]](#)
- Freshness and topical authority appear more prominent in AI retrieval pipelines, especially Perplexity, but not rigorously quantified vs traditional ranking. [\[60\]](#) [\[48\]](#) [\[49\]](#) [\[51\]](#) [\[52\]](#) [\[54\]](#) [\[75\]](#)

Commonly promoted without strong evidence:

- “You must have `llms.txt` to be cited by AI (including Google).” [\[41\]](#) [\[33\]](#) [\[40\]](#) [\[37\]](#) [\[38\]](#)
- “You must artificially chunk content for AI to understand/cite you.” [\[33\]](#) [\[37\]](#) [\[38\]](#) [\[39\]](#)
- “Adding schema will reliably increase AI citations.” [\[68\]](#) [\[63\]](#) [\[64\]](#) [\[65\]](#) [\[66\]](#) [\[67\]](#)
- “GEO/AEO are entirely new disciplines separate from SEO.” [\[32\]](#) [\[34\]](#) [\[35\]](#) [\[36\]](#) [\[31\]](#)
- “AI Overviews make traditional SEO obsolete.” [\[81\]](#) [\[82\]](#) [\[83\]](#) [\[80\]](#)

If you'd like, I can turn this into a checklist: "Do these first (verified), consider these next (correlative), ignore or de-prioritize these (hype)."

✱

Suppose a website already ranks well organically in Google.

Why might its page still not be cited in an AI-generated answer, and why might the brand itself still not be mentioned or recommended?

Separate these stages where possible:

retrieval

source selection

citation

brand mention

recommendation

Distinguish documented mechanisms from hypotheses or unknowns.

Even if a page ranks #1–3 in Google organic results, it can still be absent from AI-generated answers—or appear only as an anonymous link. That happens because AI search doesn't simply "cite the top organic result." It runs a multi-stage pipeline where a page can drop out at retrieval, source selection, citation, brand mention, or recommendation. ^[87] ^[88] ^[89] ^[90] ^[91] ^[92]

Below is what's documented vs hypothesized at each stage, and why strong organic rank alone doesn't guarantee visibility in AI answers.

1. Retrieval: "Is this page even in the candidate pool?"

Goal: Turn the user's query into a set of candidate passages from the web.

Documented mechanisms

- **Query decomposition (fan-out):** For many queries, Google's AI Overviews and other engines split the original question into multiple related sub-queries and retrieve results for each. A page that ranks for the main query may never appear in the sub-query result sets that actually feed the answer. ^[88] ^[93] ^[94] ^[87]
- **Hybrid retrieval (lexical + semantic):** Engines use both keyword matching (BM25-style) and vector/semantic similarity to fetch candidate passages. A page can be strong on classic ranking signals but weak on semantic match to the specific phrasing or intent the AI uses. ^[89] ^[90] ^[95] ^[96]
- **Crawl & index eligibility:** To be retrieved, a page must be indexed and allowed for the relevant AI crawlers. For Google, this includes being snippet-eligible and not blocked by robots.txt rules that affect AI features (e.g., disallowing Google-Extended where applicable). ^[93] ^[97]

Why a top-ranking page can fail here:

- It ranks for the *typed* query but not for the *fan-out sub-queries* the AI actually uses. ^{[87] [88] [93]}
- It's not indexed for the AI's retrieval corpus (e.g., behind a paywall without proper markup, blocked for certain bots, or not in the index used by that AI product). ^{[97] [93]}
- Its content doesn't match the semantic vector of the AI's reformulated question, even if it matches the original keywords. ^{[90] [89]}

Evidence status:

Query fan-out and hybrid retrieval are described in multiple analyses and align with how RAG systems work; Google's own documentation confirms snippet/index eligibility as a prerequisite. ^{[88] [89] [90] [93] [87]}

2. Source selection (reranking & filtering): "Which candidates survive to the answer?"

Goal: From dozens/hundreds of retrieved passages, select a small set (often 3–10) to ground the answer.

Documented / strongly supported mechanisms

- **Passage-level reranking:** Systems split documents into chunks (often ~100–200 words) and use cross-encoders or similar models to score each chunk for relevance, factual density, and answer quality. Many candidates are dropped here even if they passed retrieval. ^{[98] [99] [89] [90]}
- **Freshness & recency:** Especially for time-sensitive topics, newer or recently updated content is favored. Perplexity explicitly lists freshness as a key factor; Google's AI Overviews also show recency effects in practice. ^{[96] [100] [101] [102]}
- **Authority & trust signals:** Domain reputation, author credentials, corroboration across sources, and editorial quality are used to filter or down-rank low-trust pages. ^{[92] [103] [96] [97]}
- **Extractability / answer-shaped content:** Passages that read as self-contained answers (clear claim, data, definition) score higher for "extraction confidence." ^{[104] [90] [92] [97]}

Why a top-ranking page can fail here:

- It's relevant but **outdated** compared to fresher sources the AI prefers. ^{[100] [101] [96]}
- It's a long, narrative page where the key answer is buried or entangled with other topics, lowering extraction confidence. ^{[90] [97]}
- It lacks corroborating signals (few external mentions, weak entity presence) compared to other candidates that the system deems more trustworthy. ^{[103] [92] [97]}

Evidence status:

Reranking, freshness, authority, and extractability are consistently described in platform documentation (Perplexity) and technical explainers of RAG pipelines; Google's AI guidance emphasizes quality and people-first content, which maps to these signals. ^{[101] [89] [92] [96] [98] [100] [90]}

3. Citation: “Does this source get attached to a claim?”

Goal: When generating the answer, decide which passages to explicitly cite for specific claims.

Documented / strongly supported mechanisms

- **Grounding & attribution:** RAG systems attach citations to sentences that are directly supported by specific passages. If a passage isn’t used to ground a claim, it won’t be cited, even if it was in the candidate pool. ^{[91] [105] [106] [107] [108]}
- **Specificity & uniqueness:** Passages with unique, specific, verifiable facts (numbers, definitions, steps) are more likely to be cited than generic or opinionated text. ^{[105] [107] [92] [104]}
- **Redundancy & corroboration:** When multiple sources say the same thing, the system may cite only a subset (often those with higher trust or clearer phrasing). ^{[108] [92] [105]}

Why a top-ranking page can fail here:

- Its content is used implicitly (influences the answer) but not tied to a specific sentence that needs grounding, so no citation is attached. ^{[107] [91]}
- Its claims are generic or not uniquely stated, so the system cites other sources that express the fact more precisely or with better supporting context. ^{[92] [104] [105]}
- The system chooses to cite only a limited number of sources and prioritizes others with stronger trust or clarity signals. ^{[103] [92]}

Evidence status:

The grounding/citation attachment process is well described in RAG literature and multiple AI-search explainers; the preference for specific, verifiable facts is consistent across sources. ^{[106] [91] [104] [105] [107] [108]}

4. Brand mention: “Is the brand name actually written in the answer?”

Goal: Decide whether to name the brand in the visible text, not just link to the domain.

Documented / strongly supported mechanisms

- **Ghost citations are common:** Large-scale studies (e.g., Semrush + Kevin Indig, 2026) find that roughly 62% of AI citations are “ghost citations”: the domain is linked, but the brand name never appears in the answer text. ^{[109] [110] [111] [112] [113] [114]}
- **Citation ≠ mention:** Analyses show citation rates around ~75% but brand mention rates around ~38% for domain appearances, indicating that being cited does not guarantee being named. ^{[111] [112] [109]}

Why a top-ranking, cited page may still not get its brand mentioned:

- The system treats the domain as a source of facts but doesn’t need to name the brand to explain the claim (e.g., “According to a 2025 study...” with a link, but no brand). ^{[110] [109] [111]}

- The brand is not strongly associated as an entity in that topic area (weak entity graph, few third-party mentions), so the model doesn't surface the name as part of the narrative. ^[112]
^[110]

Evidence status:

The “ghost citation” phenomenon is supported by multiple 2026 studies across major AI engines; the mechanism (citation without naming) is observed, though the exact internal decision logic is not fully disclosed. ^[113] ^[114] ^[109] ^[110] ^[111] ^[112]

5. Recommendation: “Is this brand/page actively recommended as a choice?”

Goal: Go beyond citing a fact to explicitly recommending a product, service, or brand (e.g., “Best option: X”).

Documented / observed patterns

- **Recommendation is rarer than citation:** AI engines tend to cite sources for facts and definitions but are more cautious about explicit recommendations, especially in regulated or high-stakes categories. ^[105] ^[92]
- **Trust & corroboration matter more:** For recommendations, systems appear to favor sources with strong editorial review, comparative testing, or consensus signals (e.g., reputable review sites, official docs). ^[108] ^[92] ^[105]

Why a top-ranking page may not be recommended:

- It provides factual information but not comparative or evaluative content that supports a recommendation (e.g., no pros/cons, no “best for X” framing). ^[92] ^[105]
- The brand lacks sufficient trust signals (reviews, third-party validation, clear expertise) for the model to comfortably endorse it. ^[105] ^[108]

Evidence status:

Mostly observational and inferential from how RAG systems handle advice vs facts; platforms don't publish detailed recommendation policies, but the pattern (fewer explicit recommendations, higher trust bar) is consistent across analyses. ^[108] ^[92] ^[105]

Putting it together: why strong organic rank isn't enough

Even with a #1–3 ranking:

- **Retrieval:** The page may miss fan-out sub-queries or fail semantic match to the AI's reformulated intent. ^[93] ^[87] ^[88]
- **Source selection:** It may be outranked at the passage level by fresher, more extractable, or more authoritative chunks. ^[98] ^[90] ^[103]
- **Citation:** Its content may influence the answer but not be tied to a specific claim that gets a citation. ^[91] ^[107]
- **Brand mention:** It may be cited as a domain but remain a “ghost citation” with no brand name in the text. ^[109] ^[110] ^[111] ^[112]

- **Recommendation:** It may never be framed as a recommended option due to content type or trust signals. ^[92] ^[105]

Documented vs hypothesized

Well-documented / strongly supported:

- Query fan-out and hybrid retrieval can exclude top organic pages from the candidate pool. ^[87] ^[88] ^[93]
- Passage-level reranking on relevance, freshness, authority, and extractability filters out many candidates. ^[89] ^[96] ^[90] ^[98]
- Citations are attached to grounded claims; specificity and uniqueness increase citation likelihood. ^[107] ^[91] ^[105] ^[92]
- Ghost citations are common: ~62% of citations don't name the brand. ^[114] ^[110] ^[111] ^[112] ^[113] ^[109]

More hypothetical / partially unknown:

- Exact weighting of freshness vs authority vs extractability in Google's AI Overviews (Perplexity is more explicit; Google less so). ^[96] ^[100] ^[101]
- Precise internal rules for when a brand is named vs left anonymous in the answer text. ^[110] ^[112]
- Formal criteria for moving from "cited source" to "recommended option," especially in sensitive categories. ^[105] ^[108] ^[92]

If you'd like, I can map these stages to concrete diagnostics (e.g., how to check fan-out coverage, passage extractability, and ghost-citation rate for your domain).

✱

If you were advising an experienced SEO team for the next 90 days, what would you actually change because of AI Search?

Separate your recommendations into:

continue doing because traditional SEO remains foundational

change or add because AI-generated discovery creates a genuinely different requirement or outcome

test because the evidence is promising but not conclusive

do not treat as proven despite current industry claims

Include schema, structured data, entity optimisation, content structure, third-party mentions, llms.txt, AI crawler access, query fan-out and AI visibility measurement where relevant.

Here's a pragmatic 90-day plan for an experienced SEO team, explicitly separating what to keep, what to change, what to test, and what not to treat as proven in light of AI search (Google AI Overviews/AI Mode, ChatGPT, Gemini, Perplexity, etc.).

1. Continue doing (traditional SEO remains foundational)

These are non-negotiable and still the base layer for AI visibility.

Technical & crawl fundamentals

- **Crawlability, indexation, and site health**
 - Keep ensuring all important pages are discoverable, indexable, and free of major technical issues (redirect chains, 4xx/5xx spikes, canonical conflicts).
 - Maintain good Core Web Vitals, mobile-friendliness, HTTPS, and clean HTML. ^[120]
- **Robots.txt & sitemaps**
 - Keep a clean robots.txt and XML sitemaps; ensure Googlebot can access everything you want in Search and AI Overviews. ^[120]

Content quality & E-E-A-T

- **People-first, non-commodity content**
 - Continue investing in original, expert-led content with first-hand experience, data, and analysis. Google explicitly calls this more important than any other AI-specific tweak. ^[121] ^[120]
- **Clear structure for humans**
 - Keep logical headings, sections, and navigation that help readers; this also helps AI systems understand and extract content. ^[122] ^[120]

Structured data (for classic SEO)

- **Schema for rich results and clarity**
 - Continue using schema (Organization, Person, Article, Product, FAQPage, etc.) where it supports rich results and semantic clarity.
 - Do not remove existing schema; it still matters for traditional Search and helps define entities. ^[123] ^[121] ^[120]

Local & ecommerce fundamentals

- **Google Business Profile & Merchant Center**
 - For local businesses: keep GBP complete, accurate, and actively managed.
 - For ecommerce: maintain Merchant Center feeds, product schema, and structured product data. Google explicitly ties these to AI responses for local and shopping queries. ^[121] ^[120]

2. Change or add (genuinely different requirements because of AI search)

These changes are driven by how AI systems retrieve, select, and cite sources.

AI crawler access (explicitly allow search/retrieval bots)

What to change:

Audit and update `robots.txt` specifically for AI crawlers, distinguishing training vs search/retrieval bots.

- **Allow search/retrieval crawlers** so your pages can be cited:
 - OAI-SearchBot (ChatGPT Search)
 - Claude-SearchBot (if applicable)
 - PerplexityBot (+ Perplexity-User for live fetches)
- **Decide on training crawlers** based on your policy:
 - If you want to allow training: GPTBot, ClaudeBot, Google-Extended, CCBot, etc.
 - If you don't: explicitly Disallow them while keeping search bots allowed. ^[124] ^[125] ^[126] ^[127] ^[128]

Why:

If these bots can't access your pages, your content effectively doesn't exist in their index, regardless of how well you rank in Google. ^[125] ^[124]

90-day action:

- Week 1-2: Log analysis to identify AI bots hitting your site; confirm current `robots.txt` rules.
- Week 2-4: Implement a clear policy (e.g., "allow search, block training") and update `robots.txt` accordingly; monitor for errors. ^[126] ^[127] ^[125]

Entity optimization (make your brand and key concepts machine-readable)

What to add:

Treat entity clarity as a core objective, not a "nice to have."

- **Define your entity model**
 - List brand, products, methodologies, key people, and core topic entities you want to own. ^[129] ^[130]
- **Ship Organization and Person schema**
 - On homepage, About, and author pages: Organization, Person, sameAs links to official profiles (LinkedIn, Crunchbase, Wikipedia/Wikidata if available). ^[130] ^[129]
- **Enforce cross-platform consistency**
 - Align name, description, and URLs across directories, social bios, press mentions, and partner pages. ^[131] ^[129] ^[130]

- **Build corroborating mentions**
 - Prioritize high-quality, context-rich mentions (podcasts, industry reports, reputable blogs) that reinforce your entity definition. ^[129] ^[130]

Why:

Entity clarity affects whether AI systems recognize you as a distinct, trustworthy source and can correctly associate facts with your brand. ^[132] ^[130] ^[129]

90-day action:

- Days 1–30: Entity audit; implement/fix Organization + Person schema; align top 50 external profiles.
- Days 31–60: Target 10–20 high-value mentions that explicitly name and describe your brand/entity. ^[130] ^[129]

Content structure for retrieval and extraction (answer-first, fan-out aware)

What to change:

Shift from “long-form for dwell time” to “long-form that is easy to retrieve and quote.”

- **Answer-first formatting**
 - For each key question, lead with a 40–60 word direct answer, then expand with supporting detail. ^[133] ^[134] ^[135] ^[122]
- **Self-contained sections**
 - Write H2/H3 sections that can stand alone (roughly 200–500 words) so they can be lifted as passages without extra context. ^[135] ^[122]
- **Explicit Q&A and lists**
 - Use clear question headings and support with bullets/numbered steps; add FAQPage schema where it fits naturally. ^[134] ^[122]

Why:

AI systems retrieve and score passages, not whole pages. Clear, self-contained answers increase extraction confidence and citation odds. ^[133] ^[134] ^[135]

90-day action:

- Days 1–30: Identify top 20 informational pages; rewrite intros to be answer-first.
- Days 31–60: Restructure top 10 pages with clear H2s per sub-question; add supporting lists/tables. ^[122] ^[134]

Query fan-out optimization (cover the hidden sub-queries)

What to add:

Optimize for the *set* of sub-queries AI engines generate, not just the visible head term.

- **Map fan-out sub-queries**
 - Use AI tools (ChatGPT, Gemini) to run your core buyer prompts multiple times; capture the sub-queries they reveal in reasoning traces or follow-ups. [\[136\]](#) [\[137\]](#) [\[134\]](#)
 - Supplement with “People Also Ask,” AI Overview follow-ups, and related searches. [\[137\]](#) [\[136\]](#)
- **Cover every node**
 - For each major topic, ensure you have:
 - One comprehensive page or cluster that explicitly answers each sub-question.
 - Each sub-question as an H2 with an answer-first paragraph. [\[135\]](#) [\[137\]](#) [\[133\]](#)

Why:

AI Overviews and AI Mode expand one query into many hidden sub-queries; if you only cover the head term, you may be invisible to large parts of the fan-out. [\[138\]](#) [\[136\]](#) [\[133\]](#)

90-day action:

- Weeks 1–2: Pick 3–5 high-value prompts; run each 5+ times across models; cluster sub-queries. [\[137\]](#)
- Weeks 3–8: Audit existing content against those clusters; create or expand pages to cover missing nodes. [\[133\]](#) [\[137\]](#)

Third-party mentions & brand entity signals

What to add:

Systematically grow high-quality brand mentions as an AI visibility lever.

- **Target mention types**
 - Industry reports, roundups, podcasts, webinars, conference talks, and partner case studies where your brand is named and described consistently. [\[129\]](#) [\[130\]](#)
- **Focus on context**
 - Prefer mentions that tie your brand to specific concepts, methods, or categories (e.g., “X’s Y framework for Z”). [\[132\]](#) [\[129\]](#)

Why:

Brand mentions correlate strongly with AI citation/visibility (Ahrefs: ~3× stronger correlation than backlinks). While not proven causal, they reinforce entity recognition and trust. [\[139\]](#) [\[140\]](#) [\[129\]](#)

90-day action:

- Days 1–30: Audit current mentions; identify gaps in key topic areas.

- Days 31–90: Execute a focused outreach program for 10–20 high-quality mentions aligned with your entity model. ^[130] ^[129]

AI visibility measurement (new KPIs beyond rankings)

What to add:

Implement measurement for AI-specific outcomes.

- **Core metrics to track**
 - **Mention rate:** % of target prompts where your brand appears.
 - **Citation rate:** % of mentions that include a link to your domain.
 - **Share of voice/answer:** Your mentions vs competitors across engines.
 - **Ghost citation rate:** % of domain citations without brand name. ^[141] ^[142] ^[143]
- **Tools & process**
 - Use AI visibility tools (e.g., Mentionova, Flint, specialized GEO trackers) to run a fixed prompt set monthly across ChatGPT, Gemini, Perplexity, and Google AI Overviews/AI Mode. ^[142] ^[143] ^[141]

Why:

Traditional GSC data doesn't show AI citations or brand mentions; you need separate instrumentation to know if your GEO work is moving the needle. ^[143] ^[141] ^[142]

90-day action:

- Weeks 1–4: Choose a tool or build a manual process; define a 50–100 prompt set representing your core use cases.
- Weeks 5–12: Run monthly measurements; tie changes to content/entity work. ^[141] ^[142]

3. Test (promising but not conclusively proven)

Treat these as experiments with clear hypotheses and success metrics.

Schema tuned for AI extraction (beyond classic rich results)

What to test:

Adding or refining schema specifically to improve AI extraction and citation, not just rich results.

- Add FAQPage, HowTo, Article, QAPage on high-intent informational pages.
- Ensure author, datePublished, dateModified, and publisher are present and accurate. ^[122]

Rationale:

Observational studies show correlation between schema and AI citations, but controlled tests (e.g., Ahrefs 2026) show little to no causal uplift. Still, schema may help some engines or query types. ^[144] ^[145] ^[146] ^[147] ^[148] ^[149] ^[150]

Test design (90 days):

- Select 20–30 similar pages; add/enhance schema on half (A/B at page level).
- Track AI citation rate and mention rate over 6–8 weeks; compare deltas. [\[145\]](#) [\[147\]](#)

Freshness cadence for time-sensitive topics

What to test:

Systematic content refreshes (e.g., monthly data updates, quarterly rewrites) for topics where AI engines emphasize recency.

- Add visible “last updated” dates and changelogs for key pages.
- Introduce new data points, examples, or case studies on a regular cadence. [\[122\]](#)

Rationale:

Perplexity and other engines explicitly weight freshness; Google’s AI Overviews also show recency effects, but exact impact vs traditional ranking is not fully quantified. [\[151\]](#) [\[152\]](#) [\[153\]](#) [\[122\]](#)

Test design:

- Pick 10–15 time-sensitive pages; implement a strict refresh schedule.
- Measure changes in AI citation frequency and share of voice vs a control set. [\[152\]](#) [\[151\]](#)

Comparison and “best X for Y” content optimized for recommendation

What to test:

Creating or enhancing comparison pages designed to be used as recommendation sources.

- Add structured comparison tables, pros/cons, and clear “best for” segments.
- Cite your own data or testing methodology to boost trust. [\[122\]](#)

Rationale:

AI engines appear to favor well-structured, comparative content for recommendation-style answers, but formal criteria are not fully disclosed. [\[154\]](#) [\[155\]](#) [\[122\]](#)

Test design:

- Build 5–10 high-quality comparison pages in key categories.
- Track how often they appear in recommendation-style AI answers over 2–3 months. [\[155\]](#) [\[154\]](#)

4. Do not treat as proven (despite industry claims)

Keep these in the “hype / unproven” bucket unless you see clear, controlled results in your own environment.

llms.txt and “special AI files”

- Google’s official AI optimization guide states that `llms.txt` and similar files are **not used** by Google Search for AI Overviews or AI Mode. [\[156\]](#) [\[157\]](#) [\[123\]](#) [\[120\]](#) [\[121\]](#)
- Some non-Google tools might read them, but treating `llms.txt` as a core GEO requirement is not supported by Google’s documentation. [\[123\]](#) [\[121\]](#)

Stance:

Optional experiment at most; do not prioritize over core SEO, entity work, and content structure. [\[157\]](#) [\[156\]](#)

Artificial content “chunking” for AI

- Google explicitly says there’s **no need** to break content into tiny pieces for AI; its systems can extract relevant passages from multi-topic pages. [\[156\]](#) [\[157\]](#) [\[121\]](#) [\[123\]](#)
- Logical sections are good; micro-chunking purely for AI is not a proven requirement. [\[121\]](#) [\[123\]](#)

Stance:

Focus on clear sections and answer-first writing, not forced micro-chunks. [\[120\]](#) [\[122\]](#)

Rewriting content specifically “for AI”

- Google states that rewriting content just for AI (e.g., synonym-stuffed variants) is unnecessary; models already understand meaning and synonyms. [\[157\]](#) [\[123\]](#) [\[121\]](#)
- Optimize for clarity and answer quality, not for an imagined “AI-only” style. [\[120\]](#)

Stance:

Improve existing content for humans and extraction; don’t create parallel AI-only versions without evidence of benefit. [\[121\]](#) [\[120\]](#)

Schema as a guaranteed AI-citation booster

- Observational data shows correlation between schema and AI citations, but controlled experiments (Ahrefs 2026) found **no meaningful uplift** from adding schema alone, and even a slight dip for AI Overviews in some cases. [\[146\]](#) [\[147\]](#) [\[148\]](#) [\[149\]](#) [\[150\]](#) [\[145\]](#)
- Google’s guide says there’s no magic markup for AI. [\[123\]](#) [\[121\]](#)

Stance:

Use schema for rich results and entity clarity, but don’t promise stakeholders that schema alone will increase AI citations. [\[123\]](#) [\[120\]](#) [\[121\]](#)

“GEO/AEO are entirely new disciplines that replace SEO”

- Google frames AEO and GEO as still SEO, built on the same core systems and quality frameworks. ^[156] ^[157] ^[120] ^[121]
- The emphasis shifts (more on citations, entities, answer structure), but the foundation remains traditional SEO. ^[120] ^[121]

Stance:

Position GEO/AEO as extensions of SEO, not replacements; keep investing in core SEO while layering AI-specific tactics. ^[121] ^[120]

90-day roadmap (high level)

Days 1–30

- Audit and fix AI crawler access (robots.txt for GPTBot, OAI-SearchBot, PerplexityBot, etc.). ^[127] ^[125] ^[126]
- Entity audit; implement/fix Organization + Person schema; align top external profiles. ^[129] ^[130]
- Map query fan-out for 3–5 core prompts; identify content gaps. ^[136] ^[137]
- Set up AI visibility measurement (prompt set, tools, baseline). ^[142] ^[141]

Days 31–60

- Rewrite top 20 informational pages to be answer-first; restructure with clear H2s per sub-question. ^[134] ^[122]
- Launch focused mention-building campaign (10–20 high-quality, entity-aligned mentions). ^[130] ^[129]
- Begin schema/enrichment tests on a subset of pages; define success metrics. ^[147] ^[145]

Days 61–90

- Expand fan-out coverage: create or enhance pages for missing sub-queries. ^[137] ^[133]
- Implement freshness cadence for time-sensitive pages; add “last updated” signals. ^[122]
- Review AI visibility metrics; double down on tactics that move mention/citation rates; pause or adjust low-impact experiments. ^[141] ^[142]

If you want, I can turn this into a concrete checklist with owners, effort estimates, and example robots.txt snippets for AI crawlers.



1. <https://www.searchenginejournal.com/googles-new-ai-search-guide-calls-aeo-and-geo-still-seo/575026/>
2. <https://almcorp.com/news/good-seo-still-good-seo-ai-search-age-google-kraham/>
3. <https://www.2pointagency.com/blog/ai-search-vs-traditional-seo/>
4. <https://www.w3era.com/blog/seo/optimize-for-chatgpt-gemini-perplexity/>

5. <https://www.amivisibleonai.com/blog/ai-seo-guide-2026>
6. <https://www.zonetechify.com/blog/how-to-optimize-your-website-for-ai-search-engines-chatgpt-gemini-and-perplexity-ai>
7. <https://witscode.com/guides/ai-llm-seo>
8. <https://www.optimizegeo.ai/blog/generative-ai-seo>
9. <https://www.imarkinfotech.com/ai-seo-vs-traditional-seo-whats-changing-in-2026/>
10. <https://www.overtopseo.com/geo-vs-seo-generative-engine-optimization-differs-traditional-search/>
11. <https://www.ritnerdigital.com/blog/the-new-seo-checklist-optimizing-for-chatgpt-gemini-and-perplexity>
12. <https://www.ferventers.com/blogs/ai-seo-vs-traditional-seo>
13. <https://strapi.io/blog/generative-engine-optimization-vs-traditional-seo-guide>
14. <https://www.cognizo.ai/blog/geo-vs-traditional-seo>
15. <https://www.digitalapplied.com/blog/seo-after-ai-overviews-complete-strategy-guide-2026>
16. <https://infinitemediarresources.com/generative-engine-optimization-ai-search/geo-vs-seo/>
17. <https://mikekhorev.com/google-ai-overview>
18. <https://claritydigital.agency/blog/how-to-optimize-for-every-major-ai-search-platform-2026>
19. <https://www.digitalapplied.com/blog/seo-after-google-io-2026-90-day-playbook>
20. <https://www.stackmatix.com/blog/traditional-seo-vs-ai-search-optimization>
21. <https://llmpulse.ai/blog/geo-seo-practical-playbook/>
22. <https://www.synfinitydynamics.com/blogs/How-to-rank-in-chatgpt-gemini-perplexity>
23. <https://pbjmarketing.com/blog/ai-seo-vs-traditional-seo>
24. <https://lavio.ai/ai-search-vs-seo-what-actually-changed-in-2025-2026/>
25. <https://morningscore.io/what-distinguishes-traditional-seo-from-generative-engine-optimization-geo/>
26. <https://www.nudgenow.com/blogs/geo-vs-seo>
27. <https://www.seo-kreativ.de/en/blog/google-ai-overviews-updates-2026-en/>
28. <https://quickseo.ai/blog/google-ai-overviews-statistics-2026-60-data-points-every-seo-should-know>
29. <https://www.digitalapplied.com/blog/google-ai-overviews-surge-58-percent-queries-seo-impact>
30. <https://www.stackmatix.com/blog/google-ai-overviews-impact-seo-2026>
31. <https://ecorpit.com/google-ai-search-optimization-guide-geo-aeo-2026/>
32. <https://www.digitalapplied.com/blog/google-official-seo-docs-generative-ai-optimization-june-2026>
33. <https://connascent.com/blog/google-ai-seo-playbook-local-business-2026>
34. <https://whitebunny.com/marketing-weekly/edition-002-google-ai-search-optimisation-guide/>
35. <https://ivabot.xyz/blog/google-ai-search-optimization-guide-2026>
36. <https://yositeup.com/blog/google-ai-guide-aeo-geo-seo-2026>
37. <https://www.semrush.com/blog/google-publishes-generative-ai-search-guide/>
38. <https://meowapps.com/google-may-2026-ai-guide-seo-engine/>
39. <https://www.pinmeto.com/news/google-ai-search-optimization-guidance-2026/>
40. <https://www.getpassionfruit.com/blog/should-i-create-an-llms.txt-file-google-s-2026-guidance-explained>

41. <https://verityscore.io/en/blog/google-ai-optimization-guide-2026/>
42. <https://www.stackmatix.com/blog/google-search-central-ai-overviews-guidance>
43. <https://www.crawlora.com/blog/oai-searchbot>
44. <https://icoda.io/blog/gptbot-guide/>
45. <https://agentsurge.io/blog/correct-robots-txt-rules-for-gptbot>
46. <https://crawlreadyai.com/blog/openai-crawler-guide>
47. <https://www.ai-advisors.ai/blog/gptbot-vs-oai-searchbot>
48. <https://firstmotion.com/insights/how-perplexity-decides-which-sources-to-cite-perplexity-citation-mechanics-explained>
49. <https://singularity.digital/insights/how-perplexity-ranking-works/>
50. <https://kompozy.io/guides/perplexity-source-selection-mechanics>
51. <https://www.datastudios.org/post/how-does-perplexity-choose-and-rank-its-information-sources-algorithm-and-transparency>
52. <https://ziptie.dev/blog/what-determines-which-sites-perplexity-cites-first/>
53. <https://citegeo.ai/blog/perplexity-ranking-guide>
54. <https://www.w3era.com/blog/seo/optimize-for-chatgpt-gemini-perplexity/>
55. <https://witscode.com/guides/ai-llm-seo>
56. <https://www.ritnerdigital.com/blog/the-new-seo-checklist-optimizing-for-chatgpt-gemini-and-perplexity>
57. <https://www.digitalapplied.com/blog/seo-after-ai-overviews-complete-strategy-guide-2026>
58. <https://infinitemediarresources.com/generative-engine-optimization-ai-search/geo-vs-seo/>
59. <https://signals.sh/blog/how-to-rank-in-google-ai-overviews>
60. <https://wpseoai.com/blog/what-does-perplexity-ai-look-for-when-ranking-content/>
61. <https://ziptie.dev/blog/faq-schema-for-ai-answers/>
62. <https://sgaindex.com/news/ai-overview3-3517d27c>
63. <https://www.aeo-rankings.com/blog/does-structured-data-win-answer-schema-citation-evidence/>
64. <https://www.aeo-rankings.com/blog/schema-markup-ai-citations-data-actually-says-contradiction/>
65. <https://deepsmith.ai/blog/does-schema-markup-improve-ai-citations-evidence>
66. <https://marcodiversi.com/blog/does-schema-help-ai-citations/>
67. <https://isagentready.com/en/blog/does-schema-markup-help-ai-citations>
68. <https://heybuffy.com/blog/does-schema-markup-improve-ai-citations>
69. <https://lumengeo.co/blog/ai-search-statistics-2026>
70. <https://www.citeflow.io/blog/brand-mentions-vs-backlinks>
71. <https://machinerelations.ai/research/ai-search-citation-factors-2026>
72. <https://www.ai-visibility-software.com/blog/ai-brand-visibility-statistics-share-voice-mentions-sentiment/>
73. <https://sqmagazine.co.uk/ai-overviews-statistics/>
74. <https://parse.gl/blog/google-ai-overviews-brand-impact>
75. <https://claritydigital.agency/blog/how-to-optimize-for-every-major-ai-search-platform-2026>

76. <https://www.2pointagency.com/blog/ai-search-vs-traditional-seo/>
77. <https://www.imarkinfotech.com/ai-seo-vs-traditional-seo-whats-changing-in-2026/>
78. <https://www.searchenginejournal.com/googles-new-ai-search-guide-calls-aeo-and-geo-still-seo/575026/>
79. <https://almcorp.com/news/good-seo-still-good-seo-ai-search-age-google-kraham/>
80. <https://lavio.ai/ai-search-vs-seo-what-actually-changed-in-2025-2026/>
81. <https://www.seo-kreativ.de/en/blog/google-ai-overviews-updates-2026-en/>
82. <https://www.digitalapplied.com/blog/google-ai-overviews-surge-58-percent-queries-seo-impact>
83. <https://www.stackmatix.com/blog/google-ai-overviews-impact-seo-2026>
84. <https://developers.google.com/search/blog/2026/05/a-new-resource-for-optimizing>
85. <https://www.digitalapplied.com/blog/google-search-console-ai-overviews-mode-controls-2026-guide>
86. <https://authoritytech.io/curated/ai-brand-mention-vs-citation-gap-2026>
87. <https://contently.com/2026/06/11/google-ai-overview-citations/>
88. <https://www.searchenginejournal.com/google-ai-overview-citations-from-top-ranking-pages-drop-sharply/568637/>
89. <https://www.xseek.io/blogs/articles/how-ai-search-works-rag-vectors-and-geo-explained-2026>
90. <https://maxaeo.ai/blog/how-ai-retrieval-works/>
91. <https://neuraladx.com/how-ai-search-works-retrieval-to-citation-pipeline/>
92. <https://cite.solutions/blog/how-ai-decides-which-sources-to-cite>
93. <https://www.rankinaioverview.com/blog/how-to-rank-in-google-ai-overviews>
94. <https://digitalstrategyforce.com/journal/why-are-your-top-ranked-pages-missing-from-ai-search-answers-in-2026/>
95. <https://patrickstox.com/ai-search/how-search-works/>
96. <https://aiboost.co.uk/inside-the-retrieval-pipeline-how-generative-engines-gather-sources/>
97. <https://digitalstrategyforce.com/journal/why-isnt-my-website-appearing-in-googles-ai-overview/>
98. <https://www.linkedin.com/pulse/how-ai-search-engines-choose-sources-5-stage-retrieval-leandro-nicor-2nu6f>
99. <https://towardsdatascience.com/rag-explained-reranking-for-better-answers/>
100. <https://firstmotion.com/insights/how-perplexity-decides-which-sources-to-cite-perplexity-citation-mechanics-explained>
101. <https://ziptie.dev/blog/what-determines-which-sites-perplexity-cites-first/>
102. <https://singularity.digital/insights/how-perplexity-ranking-works/>
103. <https://authoritytech.io/blog/how-perplexity-selects-sources-algorithm-2026>
104. <https://discoveredlabs.com/blog/how-ai-systems-decide-what-to-cite-the-technical-mechanics-of-llm-content-retrieval>
105. <https://cited.md/article/how-do-ai-engines-decide-which-sources-to-trust-in-a-generative-answer>
106. <https://horusindex.com/blog/how-ai-answer-engines-decide-which-sources-to-cite/>
107. <https://gist.github.com/joshbachynski/8024b560856ea6acc2b48b0ecd3705d0>
108. <https://seerly.app/resources/engineering/llm-citation-mechanics-behind-ai-search-visibility>
109. <https://marketingagent.blog/2026/06/09/62-of-ai-citations-dont-mention-your-brand-the-ghost-citation-problem/>
110. <https://aiboost.co.uk/brand-mention-vs-citation-frequency-divergence/>

111. <https://heybuffy.com/blog/ghost-citations-cited-not-named>
112. <https://authoritytech.io/curated/ghost-citations-ai-brand-visibility-2026>
113. <https://machinerelations.ai/research/semrush-ai-visibility-index-126-million-prompts-citation-authority-2026>
114. <https://www.elmohq.com/blog/ghost-citations>
115. <https://digitalstrategyforce.com/journal/how-do-ai-search-engines-decide-which-sources-to-show-first/>
116. <https://tao-hpu.medium.com/ai-search-engines-demystified-from-embedding-to-reranking-07caeec41015>
117. <https://www.parse.gl/blog/ghost-citations-brand-strategy>
118. <https://answerlab.au/blog/how-ai-decides-citations>
119. <https://authoritytech.io/curated/ai-brand-mention-vs-citation-gap-2026>
120. <https://www.seostrategy.co.uk/llm-optimisation/google-ai-optimisation-guide-2026/>
121. <https://www.hyperfx.ai/blog/google-ai-optimization-guide-2026>
122. <https://www.superlines.io/articles/google-ai-optimization-guide-may-2026/>
123. <https://www.jdsupra.com/legalnews/google-s-ai-optimization-guide-just-9316787/>
124. <https://authoritytech.io/blog/entity-seo-knowledge-graph-optimization-ai-engines-2026>
125. <https://www.anagram.ai/blog/ai-crawlers-explained-gptbot-claudebot-perplexitybot-and-how-to-let-them-in-2026>
126. <https://getspotted.ai/blog/robots-txt-for-ai-crawlers>
127. <https://www.digitalapplied.com/blog/ai-crawler-access-control-2026-robots-llms-txt-decision-matrix>
128. <https://geo.tenten.co/en/blog/ai-crawler-user-agent-list-2026>
129. <https://www.digitalapplied.com/blog/entity-seo-knowledge-graph-optimization-guide-2026>
130. <https://jottler.co/blog/knowledge-graph-seo>
131. <https://sunilpratapsingh.com/guides/geo/knowledge-graph-optimization>
132. <https://www.geoxylia.com/blog/entity-seo>
133. <https://sanbi.ai/blog/query-fan-out-explained>
134. <https://wellows.com/blog/how-to-optimize-for-ai-query-fan-out/>
135. <https://llmranks.io/blog/how-google-ai-mode-works-query-fan-out>
136. <https://ahrefs.com/blog/query-fan-out/>
137. <https://nobori.ai/blog/query-fan-out-optimization-hidden-sub-queries-ai-citations-2026>
138. <https://astiva.ai/blog/query-fanout>
139. <https://www.citeflow.io/blog/brand-mentions-vs-backlinks>
140. <https://machinerelations.ai/research/ai-search-citation-factors-2026>
141. <https://www.flint.com/articles/best-ai-visibility-tracking-tools>
142. https://cloro.dev/blog/llm_visibility_tracking_tools/
143. <https://airankchecker.net/blog/ai-visibility-monitoring-platforms/>
144. <https://ziptie.dev/blog/faq-schema-for-ai-answers/>
145. <https://www.aeo-rankings.com/blog/schema-markup-ai-citations-data-actually-says-contradiction/>
146. <https://deepsmith.ai/blog/does-schema-markup-improve-ai-citations-evidence>
147. <https://www.aeo-rankings.com/blog/does-structured-data-win-answer-schema-citation-evidence/>

148. <https://marcodiversi.com/blog/does-schema-help-ai-citations/>
149. <https://isagentready.com/en/blog/does-schema-markup-help-ai-citations>
150. <https://heybuffy.com/blog/does-schema-markup-improve-ai-citations>
151. <https://firstmotion.com/insights/how-perplexity-decides-which-sources-to-cite-perplexity-citation-mechanics-explained>
152. <https://ziptie.dev/blog/what-determines-which-sites-perplexity-cites-first/>
153. <https://aiboost.co.uk/inside-the-retrieval-pipeline-how-generative-engines-gather-sources/>
154. <https://cited.md/article/how-do-ai-engines-decide-which-sources-to-trust-in-a-generative-answer>
155. <https://cite.solutions/blog/how-ai-decides-which-sources-to-cite>
156. <https://infinacode.com/blog/google-may-2026-ai-optimization-guide>
157. https://note.com/doerstokyo_kb/n/nd62791931584?hl=en-US
158. <https://www.margen.net/ai-seo-2026-complete-guide-geo-aeo-aio-sxo/>
159. <https://www.therankmasters.com/insights/ai-visibility/best-ai-visibility-tools-share-of-voice>
160. <https://developers.google.com/search/blog/2026/05/a-new-resource-for-optimizing>
161. <https://meowapps.com/google-may-2026-ai-guide-seo-engine/>
162. <https://verityscore.io/en/blog/google-ai-optimization-guide-2026/>
163. <https://whitebunny.com/marketing-weekly/edition-002-google-ai-search-optimisation-guide/>