AI Visibility: Be the Human Voice AI Can Find

From SEO to AI Visibility Governance

By Basil C. Puglisi, MPA. A Human-AI Collaboration.

Current as of June 2026. Some platform claims, especially Bing Webmaster Tools features and Google Search guidance, are evolving and worth reverifying before acting.

Search visibility is changing, but not in the simple way most of the market describes it.

The old question was, "How do we rank?"

That question still matters. It has not disappeared. A page still has to be crawlable, indexable, technically sound, and useful enough to be eligible for search visibility. Google's own guidance is clear on that point. AI Overviews and AI Mode do not create a separate technical doorway with special markup, special schema, llms.txt files, or AI-only formatting tricks. From Google Search's perspective, optimizing for generative AI search is still SEO (Google Search Central, 2026a).

But that is only half the story.

The new question reaches past whether a page ranks. It asks whether an organization becomes a source that AI systems can find, retrieve, trust, synthesize, cite, and accurately represent. That shift changes the work. It moves visibility from page position to source selection. It moves content strategy from keyword placement to source architecture. It moves reporting from clicks and rank tracking to citations, answer inclusion, misattribution, and verification.

That is why the better term is not simply AEO, AIO, or GEO. Each term captures part of the shift, and each also carries a risk. AEO, or Answer Engine Optimization, often sounds like featured snippets with a new label. AIO, or AI Optimization, is ambiguous because it can mean general AI optimization or Google AI Overview optimization. GEO, or Generative Engine Optimization, has stronger academic footing because it describes visibility inside generative engines that synthesize answers from sources. The operational problem, however, is broader than optimization.

The discipline that is emerging is AI Visibility Governance.

AI Visibility Governance is the practice of making an organization's knowledge findable, retrievable, citable, accurate, and accountable inside AI-mediated answer systems. It includes SEO without stopping at SEO. It includes content quality without reducing content to production volume. It includes structured information without pretending that a special schema file can make AI trust a weak source. It includes measurement, and it treats a citation as success only when the citation is accurate.

The real shift moves visibility from ranking to governed representation.

This work is written for more than the businesses trying to sell something. SEO and AI optimization are usually framed as commercial tools, a way to rank in order to convert. That framing leaves out most of the people whose knowledge the public actually needs. A teacher, a clinician, a librarian, a local official, a domain researcher, a nonprofit, an independent expert,

none of them is selling a product, and each one holds knowledge that AI systems now summarize for millions of readers. Learning how to show up in AI matters for them most of all, because the answer a machine builds is only as good as the human voices it can find.

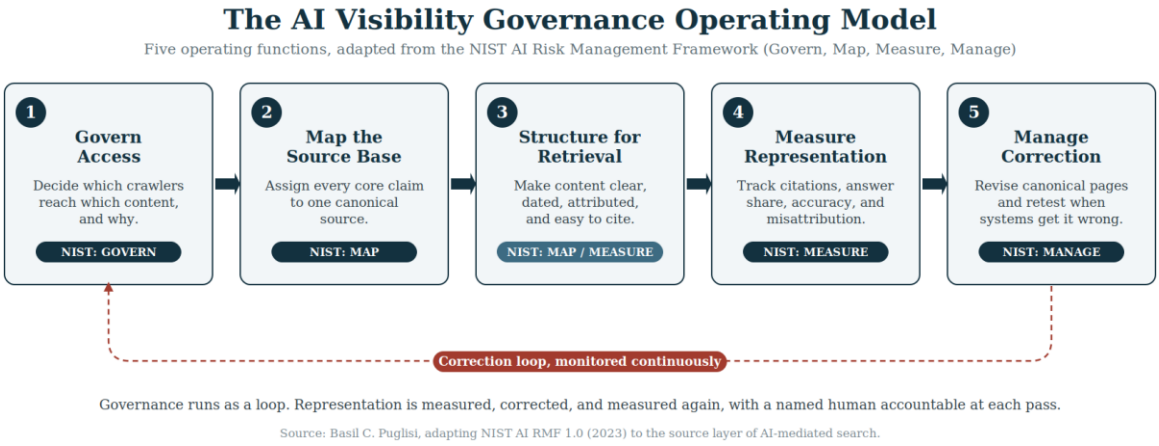

The AI Visibility Governance Operating Model.

The foundation still starts with SEO

The first mistake in the AI visibility conversation is pretending that traditional SEO no longer matters.

That is wrong.

Google's AI features still depend on the same basic visibility conditions that made search work before AI summaries entered the results page. A page has to be crawlable. It has to be indexable. It has to be eligible for snippets. It has to provide content that satisfies the user's need. Google's guidance on AI features says there are no additional technical requirements for appearing in AI Overviews or AI Mode beyond eligibility in Google Search (Google Search Central, 2025a).

It cuts through a growing market of unsupported "AI SEO" claims. Google does not use llms.txt for Search. Google does not require chunking content into special AI-friendly fragments. Google does not require special schema.org markup for AI Overviews. Google does not tell publishers to rewrite content for AI instead of people. Google's guidance says to focus on unique, non-commodity content that helps users, especially as AI search experiences invite longer, more specific, and follow-up queries (Google Search Central, 2025b; Google Search Central, 2026a).

Google has since made this position even more explicit. On June 15, 2026, the company added a dedicated subsection to its generative AI optimization guide, titled "Clarifying guidance on llms.txt files," after, in its own framing, community questions made clear that misinformation was circulating. The wording is unusually direct. Maintaining an llms.txt file is fine, and it has no effect, positive or negative, on Google Search rankings or AI Overviews, because Google Search does not use it (Google Search Central, 2026b). Earlier, on May 15, 2026, Google updated its Search spam policies to state that they apply to generative AI responses as well. Attempts to manipulate AI Overviews or AI Mode, along with the scaled content abuse, site reputation abuse, and link spam that already demote a site in traditional results, can also cost a site visibility inside Google's generative features (Google Search Central, 2026c). The governance implication is plain. The behaviors that protect traditional ranking now also protect AI citation eligibility, and the behaviors that damage one damage the other.

So the first governance rule is simple.

Do not abandon the foundation.

AI visibility begins with the same technical and editorial requirements that make content eligible for search. The organization still needs clean crawl access, clear canonical pages, internal linking, useful content, author clarity, accurate structured data where appropriate, and a reason for the source to be trusted. Without that foundation, AI optimization is just a label on tactics built on sand.

If the conversation stops there, though, it misses what has actually changed.

What changes when search becomes synthesis

Traditional search ranks documents. Generative search synthesizes answers.

That difference is not cosmetic. It changes the visibility unit. In traditional SEO, the page is the unit of competition. In AI-mediated answers, the claim, passage, source, entity, citation, and summary all become part of the visibility contest.

The GEO research gives this shift a formal frame. Generative engines gather information from multiple sources and summarize it into a response. That creates a different problem for content creators because they have less control over when and how their content appears in the generated answer. Aggarwal and colleagues describe GEO as a framework for improving visibility in generative engine responses, and they report that certain content modifications can improve visibility by up to 40 percent in their experiments. The strongest strategies include adding statistics, quotations, and citations, while keyword stuffing performs poorly or negatively (Aggarwal et al., 2024).

That finding should not be misread as permission to spam statistics or manufacture authority. The better reading is governance based. AI systems need source material that is clear, attributable, specific, and easy to synthesize without distortion. A page that says something concrete, supports it with evidence, names the source, and presents it in a clear structure is easier to retrieve and represent than a page filled with generic claims.

This reconciles the apparent conflict between Google and GEO.

Google is right that there is no special AI-only technical layer required for Google Search. GEO is also useful because it shows that content quality, evidence, and citation structure can influence how generative systems represent sources. The synthesis holds on both sides. AI visibility does not require magic markup, and it does reward governed source architecture.

From ranking metrics to citation metrics

The measurement layer is already changing, and it is changing faster than most content teams have adjusted to.

Bing's AI Performance report in Webmaster Tools made that visible first. In a February 10, 2026 public preview, Microsoft introduced reporting that helps site owners understand how their content appears in AI-generated answers across Microsoft Copilot, Bing AI summaries, and select partner integrations, including cited pages, grounding query phrases, page-level citation activity, and visibility trends over time (Microsoft Bing Blogs, 2026a). That mattered because it moved AI visibility from speculation to instrumentation.

The tooling did not stand still. At SEO Week in late April 2026, Microsoft previewed four additions to the dashboard, and these began rolling out in preview around June 16, 2026. They are Citation Share, Intents, Topics, and Compare (Microsoft Bing Blogs, 2026b). Each one closes a gap that the original preview left open. Citation Share reports the percentage of citations a site receives for a specific grounding query, so a site cited three times out of ten visible citations sees a thirty percent share. Intents classify grounding queries by purpose, such as informational, commercial, or research. Topics cluster related queries into broader subject areas. Compare overlays prior time periods so a publisher can see how citation activity moves. Microsoft has been careful to frame Citation Share as observational rather than competitive, since it does not expose competitor domains, though the metric still answers a question raw citation counts could not, which is whether a site is winning or losing the share of a given answer.

This bears directly on how the visibility dashboard should be designed, because it confirms that answer share is now a measurable concept and not a theoretical one. Google may eventually fold AI feature performance into broader Search reporting, and it has begun adding limited generative-AI metrics, yet as of this writing Bing exposes AI-specific measures, including citation share, more directly than Google Search Console does. The gap creates a practical governance lesson. The future visibility dashboard cannot stop at keyword rank, impressions, clicks, and average position. Those still matter, and they are incomplete on their own.

The new dashboard needs to ask different questions.

Which pages are cited by AI systems? Which queries trigger those citations? What share of a given answer does the organization hold? Which claims are being used as grounding material? Is the AI answer accurate? Does the citation support the claim? Are competitors cited instead? Is the organization named correctly? Is the source canonical, or is the AI citing a copied, syndicated, outdated, or weaker version?

This is where AI visibility becomes a management discipline. If a page is cited inaccurately, that is not a win. If an AI system mentions the organization but misstates the method, that is not success. If an answer uses the organization's concept without attribution, that is visibility failure hidden inside apparent influence.

The KPI layer has to mature.

AI citation count matters. Cited-page count matters. Citation share matters. Grounding-query coverage matters. So do citation precision, citation recall, misattribution rate, answer-share accuracy, correction cycle time, and source-of-truth coverage.

The click problem changes the business case

AI visibility also changes the relationship between visibility and traffic.

Pew Research Center found that Google users who encountered an AI summary clicked a traditional search result in 8 percent of visits, compared with 15 percent of visits when no AI summary appeared. Pew also found that users clicked links inside AI summaries in only 1 percent of visits (Chapekis & Lieb, 2025).

The dispute belongs alongside the finding, because governance means presenting evidence honestly, rebuttals included. Google publicly challenged the study, calling its methodology flawed and its query set unrepresentative, and the company maintains that AI features create new opportunities for people to engage with websites (Search Engine Land, 2025). The study itself drew on 68,879 searches from 900 adults over one month, of which 12,593 produced an AI summary. Critics, including former Bing executive Duane Forrester, argued that even that query

volume is too small to represent Google's overall search traffic. The finding remains directionally consistent with other click studies, and a careful reader should hold both the result and the rebuttal in view (Chapekis & Lieb, 2025).

Even granting the dispute, the strategic point holds. Citations still carry value. The value chain has changed.

In traditional SEO, rank led to click, click led to visit, and visit led to conversion or influence. In AI-mediated search, the user may get the answer without clicking. The organization may still influence the answer, and the influence may not show up as direct referral traffic. The user may remember the brand, search later, ask a follow-up question, compare options, or never leave the AI interface at all.

That creates a governance and measurement problem.

If organizations only measure clicks, they will undercount AI-mediated influence. If they only measure mentions, they will overcount weak visibility. If they only measure citations, they may miss whether the source was accurately represented.

The right model measures three things together: presence, accuracy, and downstream value.

Presence asks whether the organization appears. Accuracy asks whether it is represented correctly. Downstream value asks whether that appearance supports trust, branded search, direct traffic, qualified referral, conversion, or institutional authority.

This is a new evidence layer, and a single added column on an SEO report does not capture it.

Crawl access is now a policy decision

AI visibility is also a crawler governance problem.

OpenAI's crawler documentation makes the distinction clear. OAI-SearchBot is used to surface websites in ChatGPT search features. GPTBot is associated with crawling content that may be used to improve generative AI models. OpenAI says these settings are independent, so a webmaster can allow OAI-SearchBot while disallowing GPTBot (OpenAI, n.d.).

The distinction matters. It separates search visibility from training use.

An organization may want its public pages to appear in ChatGPT search answers while still limiting use of the same content for model training. That is a governance decision as much as a technical one. It requires legal, content, technical, and leadership alignment. The robots.txt file becomes a policy artifact.

This marks a major difference from traditional SEO governance. Historically, blocking a crawler usually meant giving up search visibility. In the AI era, crawler policy has more categories. There are search crawlers, training crawlers, user-triggered fetchers, platform-specific bots, and possible third-party retrieval systems. Each one can carry different risks and benefits. The most common and most damaging error in this area is also the quietest, a blanket block added in an earlier year to keep AI from reusing content, which then silently removes the organization from AI answers entirely.

The organization needs a crawler register.

That register should identify each known AI crawler, its stated purpose, the current access rule, the business rationale, the owner, the review cadence, and the evidence used to validate behavior. The register should also account for infrastructure beyond robots.txt, since content

delivery networks and bot-protection layers can block AI crawlers even when robots.txt allows them. Crawler policy left out of AI visibility work leaves that work incomplete. It may optimize content for systems that cannot access it, or it may expose content in ways the organization did not intend.

Being cited is not being represented

Showing up in an AI answer is only half of visibility. The other half is being represented correctly. A citation that gets your method wrong still carries your name into the answer, and it carries a version of you that you never wrote.

The most dangerous mistake in AI visibility is treating citation as proof.

Generative systems can cite sources badly. They can cite sources that do not support the claim. They can cite broken links. They can cite syndicated copies instead of original sources. They can fabricate references. They can give the appearance of research without the discipline of verification.

Liu, Zhang, and Liang evaluated generative search engines and found a serious verifiability gap. On average, only 51.5 percent of generated sentences were fully supported by citations, and only 74.5 percent of citations supported their associated sentence (Liu et al., 2023). That finding matters because it gives AI Visibility Governance two necessary metrics: citation recall and citation precision.

Citation recall asks whether the generated claims are fully supported by citations. Citation precision asks whether the citation actually supports the specific claim attached to it. A system can look credible while failing both tests.

The Tow Center for Digital Journalism found similar problems in the real world of AI search. Its study compared eight AI search engines across sixteen hundred queries and found that the tools collectively provided incorrect answers to more than 60 percent of tested news-source queries. The failures included incorrect source identification, fabricated URLs, broken links, and attribution to copied or syndicated content instead of original publishers (Jaźwińska & Chandrasekar, 2025).

This is why AI visibility needs human review.

Being cited is not enough. Being cited correctly matters. Being cited to the canonical source matters. Being summarized accurately matters. Being represented in a way that preserves the author's meaning matters.

Visibility without verification is exposure dressed as governance.

This is where verification turns from a safeguard into a visibility advantage. As AI answers get questioned more often, and as readers and platforms put more weight on whether a claim can be checked, the sources that are easy to verify are the ones that get clicked. A reader who wants to confirm what the AI said follows the citation that looks trustworthy, lands on the page that actually supports it, and stays. The same logic rewards a clear blog post, a clean website, or a well-sourced book. In a low-click world, trust is what earns the click that remains.

That is why AI Visibility Governance includes source custody. Source custody means the organization knows which source supports which claim, which page is canonical, when it was last updated, who owns it, and whether its citations have been checked. In traditional SEO, a citation helped credibility. In AI Visibility Governance, a verifiable citation becomes an auditable control and, increasingly, the reason a reader clicks through at all.

The organization needs source-of-truth architecture

Most organizations do not have an AI visibility problem first. They have a source-of-truth problem first.

They have scattered pages, outdated bios, duplicated service descriptions, inconsistent terminology, unsupported claims, weak author attribution, old PDFs, dead links, unstructured case studies, and content written for campaigns instead of retrieval. Then they wonder why AI systems misstate what they do.

AI systems do not fix organizational ambiguity. They amplify it.

The first tactical layer is a knowledge inventory. The organization should identify its core claims, services, methods, definitions, proof points, case examples, authors, and authority signals. Then it should assign each one to a canonical source.

If the organization has an original framework, it needs a canonical framework page. If it has a methodology, it needs a definition page. If it makes a claim about results, it needs evidence. If the founder or subject matter expert matters, the author entity needs a clear biography, publication history, and consistent naming. If a term is often misunderstood, the organization needs a clean explainer that says what the term means and what it does not mean.

Take the common case. Instead of three competing blog posts that each describe a service a little differently, the organization consolidates them into one canonical definition page with a clear date and a named author. When the source is unified, the AI has no conflicting versions to choose between, and the answer it builds points to the page the organization actually controls.

This is where AI visibility becomes operational.

A governed source base gives AI systems clearer material to retrieve. It gives humans a better way to audit AI outputs. It gives the organization a correction path when a system misrepresents the source. It also gives the content team a strategic filter, which is to strengthen the source layer rather than simply publish more.

A practical AI Visibility Governance model

AI Visibility Governance can be organized around five operating functions.

First, govern access. Decide which crawlers may access which content and why. Maintain a crawler register. Review robots.txt and metadata directives, and check the content delivery network and security layers for blocks that override them. Separate search visibility from training use where platforms allow that distinction.

Second, map the source base. Identify canonical pages, core claims, entity records, author profiles, framework pages, evidence pages, and update histories. Remove ambiguity where multiple pages compete to explain the same concept.

Third, structure for retrieval. Use clear headings, concise definitions, visible authorship, dates, evidence, FAQs, tables, and internal links. Structured data can help where it already fits normal search practice, and it is not a magic AI schema.

Fourth, measure representation. Track where AI systems cite, mention, summarize, omit, or misstate the organization. Use prompt sets across systems. Use first-party instrumentation where it exists, such as Bing's AI Performance report and its Citation Share, Intents, and Topics views, and supplement it with cross-platform checks, since each tool sees only its own

ecosystem. Record citation accuracy, answer share, competitor presence, grounding queries, and misattribution.

Fifth, manage correction. When systems misrepresent the source, revise the canonical page, clarify the claim, improve attribution, correct outdated pages, strengthen external corroboration, and retest. Governance operates as an ongoing correction loop, not a project that ends.

None of these five functions is self-executing. Each needs a named human owner with the authority to act and the accountability to answer for the result: a technical and legal owner for access, a knowledge lead for the source base, a content owner for retrieval structure, an analytics or governance lead for measurement, and a named editor with the authority to change canonical pages for correction. A function with no owner is a process the machine can run and no person can be held to. That single requirement, a named human bound to each decision, is what separates governance from a checklist.

This structure aligns naturally with the NIST AI Risk Management Framework, which organizes AI risk work through Govern, Map, Measure, and Manage functions (NIST, 2023). AI Visibility Governance adapts that logic to the source layer. The risk runs deeper than an AI system giving a wrong answer. The deeper risk is that the organization has no governed way to know, correct, or measure the answer.

What this means for experts, small firms, and institutions

Large brands often have authority signals by default. They have backlinks, press coverage, Wikipedia pages, institutional mentions, and years of digital footprint. Independent experts, consultants, local institutions, and smaller firms have to be more deliberate.

The opportunity is not to chase every broad keyword. The opportunity is to become the canonical source for a narrow, valuable concept.

For an independent expert, that means naming the method clearly, defining it in one authoritative place, supporting claims with evidence, publishing examples, aligning profiles across platforms, and making the work easy for AI systems to retrieve and explain. For a service business, it means explaining the service in language that answers real buyer questions, not marketing slogans. For a nonprofit, it means making programs, eligibility rules, outcomes, and local authority signals clear and current. For a research organization, it means maintaining citation hygiene, version control, and human review before AI-assisted work enters the public record.

None of this requires a budget to begin. The minimum first step is to pick the single most important thing the organization knows, the one claim, definition, or eligibility rule people most need, and put it on one clear, dated, authored page that AI systems can crawl. One canonical source done well is worth more in an AI answer than a hundred scattered posts.

The tactic is to become a better source.

What should not be done

Some practices deserve an explicit warning.

Do not treat AEO, AIO, or GEO as magic labels. Do not buy unsupported claims about special AI schema. Do not create llms.txt files for Google Search visibility under the belief that Google uses them for that purpose, a belief Google has now formally and repeatedly corrected. Do not stuff pages with keywords, because traditional keyword stuffing already degraded quality and GEO

experiments show it does not solve generative visibility. Do not create generic AI-written summaries that add no original value. Do not manufacture inauthentic mentions. Do not confuse being cited with being verified.

Most important, do not let the machine become the only auditor of how the machine represents you.

The human role does not disappear. It becomes more important.

AI systems can retrieve and synthesize. They can compare sources. They can compress complexity. They do not own the organization's claims, authority, risk, or reputation. A named human or accountable team still has to decide what is true, what is canonical, what is safe to publish, what is worth correcting, and what should be rejected.

That is the governance line.

Why this matters even if you are not selling anything

Most writing about search optimization assumes a commercial motive: rank in order to sell, earn a citation in order to convert. That motive is real, and it is also narrow. The people who hold the most reliable knowledge in any field are often the ones with the least commercial reason to optimize for it.

When those voices stay absent, the AI answer does not go quiet. It gets built from whatever is present and easy to retrieve, which skews toward marketing content, search-optimized filler, and a rising share of other AI-generated text. The pool the machine draws from narrows toward the sources with a budget instead of the sources with the knowledge.

There is evidence, by analogy, that representation pools can run narrow. In a 2023 Harvard study, researchers compared an early GPT model against World Values Survey responses from 94,278 people across 65 nations. The model aligned most closely with the United States and culturally similar nations, and its resemblance to human responses fell as cultural distance from the United States grew, an inverse correlation of about negative 0.70 (Atari et al., 2023). That study measured cultural alignment, not commercial dominance, so it stands as adjacent evidence and not direct proof. The principle it illustrates still holds. A system trained and grounded on whoever shows up will represent whoever shows up, and underrepresent everyone who does not.

This is where the non-commercial voice carries weight that no brand can buy. When a clinician publishes a clear, sourced explanation of a treatment, the system has something true to retrieve instead of something invented or something sold. When a researcher names a method in one authoritative place, the answer can cite the method correctly instead of paraphrasing a weaker secondhand version. When a nonprofit states its program rules plainly, a person asking an AI about eligibility gets the real answer instead of a guess. Each clear human source improves the answer for every reader who never sees the source at all.

The citation pollution problem makes the stakes sharper. As fabricated and AI-generated references enter the public record, the systems that absorb that record inherit its errors. A web filled with confident, sourceless, machine-written summaries trains and grounds the next round of machine-written summaries, and the human signal thins with each pass. The countermeasure runs past the technical. More real human voices belong in the material, published clearly enough that AI systems can find them, cite them, and represent them without distortion.

This is the Augmented Intelligence case stated at the level of the public record. Human knowledge and machine synthesis together produce a better answer than synthesis drawing on a

thinning, commercialized, self-referential pool. For the people who are not selling anything, showing up is how human knowledge stays inside the answers that machines now give on its behalf.

Conclusion

The future of search is a source governance problem before it is an SEO tactic.

Traditional SEO still matters, because AI visibility begins with crawlability, indexing, technical clarity, and useful content. What changed sits on top of that foundation. Generative systems now synthesize answers from many sources, the measurement layer has shifted from clicks to citations and answer share, and a citation is worth nothing until it is accurate. Governance is what holds those pieces together.

The conclusion is that SEO is no longer enough by itself, even though it is far from dead.

Organizations now need AI Visibility Governance, a disciplined operating model for becoming findable, retrievable, citable, accurate, and accountable inside AI-mediated answers.

The winners will be the organizations that become the clearest source. Volume alone will not earn the citation.

This extends past business. Every clear human voice that shows up, commercial or not, makes the shared answer better for everyone who relies on it.

Disclaimer

References

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). GEO: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2024)* (pp. 5–16). Association for Computing Machinery. arXiv. https://arxiv.org/abs/2311.09735

Atari, M., Xue, M. J., Park, P. S., Blasi, D. E., & Henrich, J. (2023). Which humans? PsyArXiv. https://scholar.harvard.edu/files/henrich/files/which_humans_09222023.pdf

Chapekis, A., & Lieb, A. (2025, July 22). Google users are less likely to click on links when an AI summary appears in the results. Pew Research Center. https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

Google Search Central. (2025a, December 10). AI features and your website. Google for Developers. https://developers.google.com/search/docs/appearance/ai-features

Google Search Central. (2025b, May 21). Top ways to ensure your content performs well in Google's AI experiences on Search. Google for Developers. https://developers.google.com/search/blog/2025/05/succeeding-in-ai-search

Google Search Central. (2026a, May 15; updated June 15, 2026). Optimizing your website for generative AI features on Google Search. Google for Developers. https://developers.google.com/search/docs/fundamentals/ai-optimization-guide

Google Search Central. (2026b, June 15). Clarifying guidance on llms.txt files [subsection]. In *Optimizing your website for generative AI features on Google Search*. Google for Developers. https://developers.google.com/search/docs/fundamentals/ai-optimization-guide

Google Search Central. (2026c, May 15). Spam policies for Google web search. Google for Developers. https://developers.google.com/search/docs/essentials/spam-policies

Jaźwińska, K., & Chandrasekar, A. (2025, March 6). AI search has a citation problem. Columbia Journalism Review. https://www.cjr.org/tow_center/we-compared-eight-ai-search-engines-theyre-all-bad-at-citing-news.php

Liu, N. F., Zhang, T., & Liang, P. (2023). Evaluating verifiability in generative search engines. In *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 7001–7025). Association for Computational Linguistics. arXiv. https://arxiv.org/abs/2304.09848

Microsoft Bing Blogs. (2026a, February 10). Introducing AI Performance in Bing Webmaster Tools public preview. https://blogs.bing.com/webmaster/February-2026/Introducing-AI-Performance-in-Bing-Webmaster-Tools-Public-Preview

Microsoft Bing Blogs. (2026b, June 16). New AI visibility insights in Bing Webmaster Tools: Intents, Topics, Citation Share, Compare. https://blogs.bing.com/search/June-2026/New-AI-Visibility-Insights-in-Bing-Webmaster-Tools-Intents-Topics-Citation-Share-Compare

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce. https://www.nist.gov/itl/ai-risk-management-framework

OpenAI. (n.d.). Overview of OpenAI crawlers. OpenAI Developers. https://developers.openai.com/api/docs/bots

Search Engine Land. (2025, July 24). Google AI Overviews are hurting clicks: Pew study. https://searchengineland.com/google-ai-overviews-hurting-clicks-study-459434

#AIassisted using the HAIA Ecosystem