XenGrowth

AI SEARCH & DISCOVERY

SEO, GEO, AEO and LLMO: What Actually Changes?

AI-search acronyms overlap more than most service menus admit. Here is what the labels mean, what changes in practice, and which work remains useful across discovery systems.

An implementation guide · Reviewed August 26, 2026 · 12 min read
xengrowth.com/resources/seo-geo-aeo-llmo-ai-search-guide

Contents

01

Begin with the part the acronyms share

The labels arrived faster than the consensus. SEO usually means search engine optimization. AEO is commonly used for answer engine optimization. GEO refers to generative engine optimization, while LLMO, AISO, AIO, and GSO are competing ways to describe visibility in model-mediated or generated search. SXO keeps the visitor experience in view. In this guide, LSO means local search optimization. None of these names creates a separate law of retrieval.

Google's current AI-search guidance places the work inside established SEO: make pages crawlable, publish useful and non-commodity information, support claims, use meaningful media, and give visitors a good page experience. Bing's guidance follows a similar line for grounding in AI answers. A company does not need eight disconnected optimization programs. It needs one sound discovery system and a few additional ways to test visibility.

Label	Useful working definition	Durable focus
SEO	Organic search optimization and the technical foundation	Crawlability, relevance, authority, page experience
AEO	Answer engine optimization	Direct answers, evidence, clear passages
GEO / GSO	Visibility and citation in generated search	Source-worthiness, entities, retrieval testing
LLMO / AISO / AIO	Broader labels for model-mediated discovery	The combined SEO, answer, evidence, and measurement system
SXO	Search experience optimization	Intent match, clarity, accessibility, conversion quality
LSO	Local search optimization in this guide	Local entities, relevance, reputation, geographic intent

02

Understand how generated search expands a question

A generated answer may not rely on one literal query. Google describes a query fan-out process that issues related searches across subtopics and data sources. Retrieval-augmented generation can then bring selected passages into the context used to compose an answer. This is one reason a clear topic structure is more useful than repeating the same phrase across many URLs.

Build a central source page for the decision, then connect genuinely distinct support: definitions, implementation steps, comparisons, examples, constraints, measurement, and failure modes. Use descriptive internal links so a reader—and a crawler—can understand why the next page matters. One

topic per URL does not mean one wording variation per URL.

> *Use descriptive internal links so a reader—and a crawler—can understand why the next page matters.*

03

Make the evidence easy to retrieve and hard to misunderstand

Answer the main question early, then earn the nuance. Clear headings, short definitions, tables, lists, and visible source links help people scan and help systems locate relevant passages. The page still needs substance: first-hand experience, original data, a useful synthesis, or a concrete operating method. Formatting cannot rescue commodity copy.

Separate facts from interpretation. Give dates to time-sensitive statements. Identify the organization, product, place, or person consistently. When a claim depends on a study, link the study; when it comes from your own measurement, explain the sample and method. Accurate structured data can reinforce those visible relationships, but markup should never describe proof the page does not show.

- Direct answer followed by necessary context
- Named primary sources close to material claims
- Original examples with scope and limitations
- Stable entity names, authorship, dates, and ownership
- Images and diagrams that explain—not decorate—the subject

> *When a claim depends on a study, link the study; when it comes from your own measurement, explain the sample and method.*

04

Treat GEO studies as evidence, not a promise

The widely discussed GEO benchmark paper reported visibility gains of up to 40 percent for some optimization methods in its experimental setting. “Up to” and “in this benchmark” are important. The paper also found that the effect varied by domain, and later empirical work has documented meaningful differences between engines, languages, query wording, freshness, source selection, and interface design.

The responsible use of this research is to form hypotheses: cite sources precisely, improve factual fluency, add relevant statistics when they clarify the point, and test whether your pages appear for representative questions. It is not responsible to sell a universal 40 percent uplift. Generated-search ecosystems change too quickly, and citations do not automatically equal authority, rankings, visits, or clients.

> *Generated-search ecosystems change too quickly, and citations do not automatically equal authority, rankings, visits, or clients.*

05

Avoid the scaled-content trap

Prompt research is useful when it reveals a different job. It becomes harmful when every phrasing receives a thin page whose only purpose is to capture a variation. Google explicitly warns that creating many pages for query variations primarily to manipulate search or AI features can be scaled content abuse. Bing similarly warns against automated content at scale and keyword manipulation.

Before approving a page, ask what new decision it helps someone make. A comparison, calculator, annotated workflow, implementation checklist, evidence review, or industry-specific constraint can deserve its own URL. A synonym swap cannot. This standard is how a large editorial library remains useful instead of becoming a content landfill.

06

Measure discovery without inventing certainty

Track index coverage, search impressions, query clusters, citations, referenced URLs, AI referrals where available, branded demand, engaged visits, and assisted conversions. Bing Webmaster Tools can expose cited pages and grounding queries in its AI Performance reporting, but Bing is clear that citation count is not a ranking, authority score, or direct measure of clicks.

Build a representative question set across informational, commercial, navigational, research, creation, problem-solving, and local intent. Test it periodically across the discovery surfaces that matter to your audience. Then connect diagnostic visibility with qualified conversations, pipeline, and revenue. The final measure is not whether an answer engine noticed the page. It is whether the right person found enough clarity to take a useful next step.

Key Takeaways

- SEO remains the technical and editorial foundation for AI Overviews, AI Mode, answer engines, and other generated discovery experiences.
- GEO, AEO, LLMO, AISO, AIO, and GSO are useful planning labels, not separate channels with universally agreed definitions.
- The durable work is crawlability, explicit answers, original evidence, consistent entities, trustworthy references, useful media, and a strong visitor experience.
- Do not build a near-duplicate page for every prompt variation; group related questions around a distinct reader decision or task.
- Measure citations and visibility as diagnostics, then follow the journey through qualified demand and revenue.

Sources

1. Google Search Central: AI features and your website — https://developers.google.com/search/docs/fundamentals/ai-optimization-guide
2. Google Search Central: Creating helpful, reliable, people-first content — https://developers.google.com/search/docs/fundamentals/creating-helpful-content
3. Bing Webmaster Tools: AI Performance — https://www.bing.com/webmasters/help/ai-performance-9f8e7d6c
4. Bing Webmaster Guidelines — https://www.bing.com/webmasters/help/webmaster-guidelines-30fba23a
5. Aggarwal et al.: GEO—Generative Engine Optimization — https://arxiv.org/abs/2311.09735
6. Empirical analysis of generative search engines — https://arxiv.org/abs/2509.08919
7. Empirical study of Chinese generative search engines — https://arxiv.org/abs/2607.15771