

GEO & AEO

Winning Visibility in AI Answers: ChatGPT, Perplexity, and Google AI Overviews

A SEOWing whitepaper by GOLDENWING Creative Studios

Published [2026] · seowing.io · Version 1.0

Executive Summary

For two decades, the goal of SEO was a ranking: get your page into the ten blue links, ideally near the top, and earn the click. That goal is no longer sufficient. A growing share of searches now end not with a list of links but with a **synthesised answer** — generated by ChatGPT, Perplexity, Google's AI Overviews, and a widening field of answer engines. The user reads the answer. Whether they ever click through to your site is now a secondary event, and increasingly an optional one.

This changes the objective. It is no longer enough to *rank*; you need to be **cited** — to be one of the sources the AI reads, trusts, and names when it composes its answer. Two overlapping disciplines have emerged to pursue this:

- **GEO — Generative Engine Optimisation:** optimising your content and site so that generative AI systems select, quote, and attribute you as a source.
- **AEO — Answer Engine Optimisation:** optimising so your content is chosen as *the* answer to direct questions, in AI answers and in classic answer surfaces like featured snippets and voice.

The two are close cousins and are often pursued together; this paper treats them as a combined discipline for winning visibility in AI-mediated search.

The mechanics that drive GEO/AEO are not mystical. Answer engines favour content that is **clearly structured, factually grounded, machine-readable, and demonstrably trustworthy**. That translates into concrete, buildable signals: clean structured data (schema.org / JSON-LD), an `llms.txt` file that guides AI crawlers, strong **E-E-A-T** (Experience, Expertise, Authoritativeness, Trustworthiness), well-defined **entities**, and content organised to answer real questions directly.

This whitepaper explains how AI answer engines choose their sources, details the signals that influence that choice, provides a practical GEO/AEO checklist you can act on this week, and covers the hardest new problem: **how to measure visibility in AI answers** when there is no ranking report to read. Throughout, we show how SEOWing operationalises these signals — schema generation, `llms.txt`, E-E-A-T and entity configuration — so the work is a matter of running the right tools, not hand-editing JSON.

The strategic message: the sites that get cited in AI answers in 2026 are not the ones that gamed a ranking. They are the ones that made themselves the **easiest, clearest, most trustworthy source to quote**. GEO and AEO are how you engineer that on purpose.

1. The Shift: From Ranking to Being Cited

1.1 What actually changed

Classic search is an *index-and-rank* system: the engine crawls pages, indexes them, and returns an ordered list. The user does the synthesis — they read several results and form a conclusion.

Answer engines invert this. They perform *retrieve-and-synthesise*: the system retrieves candidate sources, reads them, and composes a single answer, often citing a handful of sources inline. The synthesis that the user used to do in their head is now done by the model. This has three consequences that reshape the SEO objective:

1. **The click becomes optional.** Many queries are answered in full on the answer surface. "Zero-click" is no longer an edge case; for informational queries it is becoming the default. [Quelle einfügen — als illustrativ/Schätzung behandeln]
2. **Citation replaces ranking as the prize.** Being *named* as a source in the answer — with a link and attribution — is the new position one. Visibility now means "quoted," not just "listed."
3. **The unit of competition shrinks.** Engines often quote a *sentence or a passage*, not a whole page. Your job is to make individual claims cleanly extractable, not just to make a page broadly relevant.

1.2 Why classic SEO is necessary but not sufficient

None of this makes traditional SEO obsolete. Answer engines still lean heavily on the crawlable, indexable web; a page that cannot be crawled or understood cannot be cited. Technical health, good metadata, internal linking, and page quality remain table stakes — they are the price of admission.

But they are no longer the whole game. A page can rank respectably in classic results and *still* never be quoted by an answer engine because its claims are buried in prose, its trustworthiness is invisible to a machine, and its structure gives the model nothing clean to lift. GEO/AEO is the layer that turns a rankable page into a **quotable** one.

2. How AI Answer Engines Choose Their Sources

You cannot optimise for a black box you do not understand. While each engine is proprietary and evolving, their source-selection behaviour is broadly observable and converges on a common logic. Understanding it is the foundation of everything that follows.

2.1 The retrieval-and-synthesis pipeline

Most answer engines follow a comparable sequence:

1. **Interpret the query** — identify the user's intent and the entities involved.
2. **Retrieve candidates** — pull a set of documents from a search index and/or live web results.

3. **Read and rank passages** — assess which passages best, and most reliably, answer the query.
4. **Synthesise** — compose an answer, weaving in the strongest passages.
5. **Attribute** — cite the sources judged authoritative and directly relevant.

Your content is competing at steps 3–5: it must be *retrievable*, its passages must be *extractable and reliable*, and the source must be *judged trustworthy enough to name*.

2.2 The five properties engines reward

Across engines, five properties consistently increase the odds of being selected and cited:

- **Clarity and structure.** Content that states claims directly, uses clear headings, and answers questions in self-contained passages is far easier for a model to lift and quote than the same information buried in a meandering narrative.
- **Machine-readability.** Structured data (schema.org / JSON-LD) tells the engine *what* your content is — an article, a product, an FAQ, an organisation — and disambiguates the entities on the page. It converts prose into facts a machine can rely on.
- **Factual grounding and specificity.** Concrete, verifiable, specific statements — with numbers, dates, and named entities — are more quotable than vague generalities. Engines prefer sources they can trust not to be hallucinating.
- **Trust and authority (E-E-A-T).** Signals that a real, qualified, accountable entity stands behind the content — identified authors with credentials, an organisation with a verifiable identity, citations, evidence of first-hand experience — raise the probability of being cited, because answer engines inherit search's strong bias toward trustworthy sources.
- **Freshness and consistency.** Current, internally consistent, and well-maintained content is favoured over stale or contradictory pages, especially for time-sensitive queries.

2.3 The uncomfortable implication

Notice what these properties have in common: **they are all things you can engineer**. None of them require insider access or tricks. They require making your content clearer, more structured, more specific, and more visibly trustworthy — which is, not coincidentally, also good for human readers. GEO/AEO rewards genuine quality made machine-legible. That is the whole game.

3. The Building Blocks of AI-Answer Visibility

This section turns the five properties into concrete, buildable signals — and shows how SEOWing produces each without hand-editing code.

3.1 Structured data (schema.org / JSON-LD)

Structured data is the single highest-leverage GEO/AEO investment because it does the model's disambiguation work for it. Instead of *inferring* that a passage is a question-and-answer pair, an author bio, or a product with a price, the engine is *told*, unambiguously, in JSON-LD.

Priority schema types for answer-engine visibility:

Schema type	What it enables	Best for
FAQPage	Direct question→answer extraction	Any page answering common questions
Article / BlogPosting	Authorship, date, publisher attribution	Editorial and blog content
Organization	A verifiable entity identity for your brand	Site-wide identity, sameAs links
Person	Author identity and credentials	E-E-A-T; author bios
Product / Offer	Structured product facts	E-commerce, comparisons
HowTo	Step-by-step extractable instructions	Procedural queries
LocalBusiness	Location, hours, service area	Local and "near me" answers
BreadcrumbList	Site structure and context	Navigation and hierarchy

A critical, often-missed point: **your structured data must be a single, clean, non-conflicting graph**. Multiple plugins each emitting their own partial schema produces contradictions that undermine trust. SEOWing is designed to be the *sole* source of structured data on the site, emitting one coherent JSON-LD graph — Organization, LocalBusiness, Service, FAQ, Breadcrumb — with no plugin conflicts. In practice you generate and attach schema with `seo_generate_schema` and `seo_set_page_schema`, or let autopilot apply the recommended set.

3.2 llms.txt — guiding the AI crawlers

Just as `robots.txt` guides search crawlers, `llms.txt` is an emerging convention for guiding AI systems: a structured, machine-readable file at your site root that tells LLM-based crawlers what your site is about, which content matters, and how it is organised. It is a direct, low-effort signal to the exact systems you are trying to be cited by.

SEOWing generates and publishes this for you with `seo_generate_llms_txt` — pulling in your key pages, posts, and a curated description — and writes it to the site root. It is one of the cleanest examples of a GEO-native signal that simply did not exist in the classic-SEO toolkit.

3.3 E-E-A-T: making trust machine-visible

Experience, Expertise, Authoritativeness, Trustworthiness is Google's long-standing quality framework, and answer engines inherit its logic: they preferentially cite sources that appear genuinely credible. The GEO/AEO challenge is that E-E-A-T is often *implicit* to human readers but *invisible* to machines. Your firm may have decades of experience and top credentials — but if none of that is expressed in a structured, extractable way, the engine cannot weigh it.

Making trust machine-visible means expressing, concretely:

- **Experience** — evidence of first-hand, real-world involvement (case studies, original data, "we tested this").
- **Expertise** — identified authors with real, stated credentials and specialisations.
- **Authoritativeness** — the organisation's identity, accreditations, certifications, and years in business.
- **Trustworthiness** — accurate contact and business information, transparent ownership, and citations to reputable sources.

SEOWing's `seo_set_eeat` captures accreditations, certifications, expert authors, founding date, years in business, and specialisations as structured signals, and `seo_set_local` adds verifiable local-business details — turning your real credibility into signals an answer engine can actually read.

3.4 Entities: being an unambiguous "thing"

Modern search and answer engines think in **entities** — distinct, disambiguated things (a specific company, person, place, product) — not just keyword strings. To be reliably cited, your brand, your authors, and your key subjects need to be recognisable as consistent entities across the web.

Practical entity work:

- Maintain a consistent `Organization` identity in schema, with `sameAs` links to authoritative profiles (official social, industry registries, knowledge bases).
- Give authors consistent, structured `Person` identities linked to their credentials.
- Use precise, consistent naming for your products, services, and concepts so the engine can bind mentions to one entity rather than scattering them.

Entity consistency is what lets an engine confidently say "according to [your brand]" — because it knows, unambiguously, who that is.

3.5 Answer-shaped content

Finally, the content itself must be *shaped* to be answerable. This is the most editorial of the building blocks and the one no plugin can fully do for you — though the right tooling makes it far faster:

- **Lead with the answer.** State the direct answer first, then elaborate. Engines lift the concise claim, not the wind-up.
- **Use real questions as headings.** Mirror how people actually ask, so a passage maps cleanly to a query.
- **Write self-contained passages.** Each section should make sense lifted out of context, because that is exactly how it will be quoted.
- **Be specific and cite.** Numbers, dates, named sources. Specificity signals reliability.

SEOWing supports this loop with `seo_analyze_page` and `seo_get_suggestions` to surface where content is thin or unstructured, and `seo_get_playbook` for current GEO/AEO/E-E-A-T guidance tailored to your site type.

4. The Practical GEO/AEO Checklist

Use this as an operational checklist. Each item maps to a signal above; the SEOWing column shows how to execute it.

#	Action	Why it matters	How in SEOWing
1	Emit one clean JSON-LD graph; remove conflicting schema plugins	Contradictory schema erodes trust	Be sole schema source; <code>seo_generate_schema</code>
2	Add <code>FAQPage</code> schema to pages answering common questions	Enables direct Q→A extraction	<code>seo_set_page_schema</code>
3	Add <code>Article</code> / <code>Person</code> schema with real author credentials	Machine-visible E-E-A-T	<code>seo_set_page_schema</code> , <code>seo_set_eeat</code>
4	Publish an <code>llms.txt</code> at your site root	Directly guides AI crawlers	<code>seo_generate_llms_txt</code>
5	Configure a consistent <code>Organization</code> entity with <code>sameAs</code>	Unambiguous brand identity	Company settings + schema
6	Fill in accreditations, certifications, years in business	Concrete authority signals	<code>seo_set_eeat</code>
7	Add <code>LocalBusiness</code> details where relevant	Wins local/"near me" answers	<code>seo_set_local</code>
8	Rewrite key pages to lead with the direct answer	Makes passages quotable	<code>seo_get_suggestions</code> , <code>seo_analyze_page</code>
9	Use real user questions as H2/H3 headings	Maps passages to queries	Content edit + audit
10	Keep content fresh and internally consistent	Freshness is a ranking/citation factor	Bulk audit + scores
11	Ensure the site is crawlable and technically healthy	No crawl, no citation	<code>seo_full_site_audit</code>
12	Submit changes for fast indexing	Speeds discovery	<code>seo_index_now</code> , sitemaps

Actionable sequencing. If you do only three things this quarter: (1) consolidate to a single clean schema graph, (2) publish `llms.txt` , and (3) make your E-E-A-T machine-visible with structured author and organisation data. Those three deliver the largest share of GEO/AEO upside for the least effort, and SEOWing can apply all three through autopilot with a dry-run preview first.

5. Measuring Visibility in AI Answers

The hardest part of GEO/AEO is not doing the work — it is *proving it worked*. Classic SEO had rank trackers. AI-answer visibility has no single equivalent, because answers are personalised, non-deterministic, and vary by engine, phrasing, and session. Measurement is genuinely harder. Here is how to approach it pragmatically.

5.1 Why it is hard

- **Non-determinism.** Ask the same question twice and the answer — and its cited sources — may differ.
- **No universal rank report.** There is no one dashboard that shows your "position" across ChatGPT, Perplexity, and AI Overviews.
- **Attribution gaps.** When an engine answers without a click, your analytics may never see the visit, even though you influenced the user.
- **Fragmentation.** Each engine behaves differently and changes frequently.

5.2 A practical measurement framework

Rather than chasing a single metric, triangulate across four layers:

Layer	What you measure	How
Citation presence	Does the engine cite you for target questions?	Prompt-based testing across engines; log which sources are named
Share of answer voice	How often you are cited vs. competitors for a query set	Track a fixed set of priority questions over time
Referral signals	Traffic and conversions from AI-answer surfaces	Analytics segments for AI referrers where identifiable
Foundational health	Are the enabling signals in place?	Schema coverage, <code>llms.txt</code> presence, E-E-A-T completeness, SEO scores

5.3 Prompt-based testing: the core method

The most direct measurement is **structured prompt testing**. Define a fixed set of the questions you want to be cited for. On a regular cadence, ask each answer engine those questions and record: Are you cited? Which competitors are cited? What passage did they quote? This gives you a repeatable, comparable "share of answer voice" over time — the closest GEO/AEO analogue to rank tracking. Standardise the prompts, the cadence, and the logging so results are comparable across periods, and accept that any single run is noisy — the *trend* across many prompts is the signal.

5.4 Leading indicators you fully control

Because outcome metrics are noisy, lean on **leading indicators** — the foundational signals that *cause* citation and that you can measure precisely on your own site:

- **Schema coverage** — what percentage of key pages carry correct, conflict-free structured data.

- `llms.txt` present and current.
- **E-E-A-T completeness** — authors, credentials, organisation identity, accreditations all populated.
- **Per-page SEO scores** — SEOWing's per-page score gives a concrete, trackable quality number across the site.
- **Technical health** — crawlability, no blocking errors, fast indexing.

These do not *prove* citation, but they are the controllable inputs. Improving them reliably raises your odds, and unlike answer-engine outputs, they move when you act and stay moved. Track them relentlessly; treat prompt-based citation testing as the outcome check on top.

5.5 The measurement takeaway

Do not wait for a perfect AI-visibility metric — it does not exist yet, and may not for a while. Instead: **instrument the leading indicators you control, run disciplined prompt-based citation testing on a fixed question set, and watch the trend.** That combination gives you an honest, actionable read on GEO/AEO performance today, without pretending the noise is not there.

6. Conclusion

Search is no longer only about being *found* — it is about being *quoted*. As ChatGPT, Perplexity, Google AI Overviews and their successors mediate more of the journey, the winners are the sites that made themselves the clearest, most structured, most trustworthy source for an engine to lift and attribute.

The good news is that GEO and AEO reward substance made legible. Clean structured data, `llms.txt`, machine-visible E-E-A-T, consistent entities, and answer-shaped content are not tricks — they are genuine quality expressed in a form machines can rely on. Every one of them is buildable, and with an AI-first plugin, most of them are a tool call away rather than a hand-edited file.

The organisations that treat AI-answer visibility as a deliberate discipline in 2026 — engineered, measured, and maintained — will compound an advantage over those still optimising only for ten blue links. The answer surface is where attention is moving. Be the source it cites.

Take the next step

SEOWing is the AI-first WordPress SEO plugin built for the answer-engine era: one clean schema graph, one-click `llms.txt`, structured E-E-A-T and entity signals, GEO/AEO best-practice playbooks, and a per-page SEO score to track it all — driven by hand or by an AI assistant over MCP. It replaces Yoast, RankMath and Imagify with a single plugin.

See how it works, or request a licence, at seowing.io. Questions? Contact us at deni@goldenwing.at — GOLDENWING Creative Studios e.U., Vienna.

Zusammenfassung (DE)

Zwei Jahrzehnte lang war das Ziel von SEO ein Ranking: unter die zehn blauen Links, möglichst weit oben. Das genügt nicht mehr. Immer mehr Suchen enden mit einer **KI-Antwort** — von ChatGPT, Perplexity oder Google AI Overviews — statt mit einer Linkliste. Das neue Ziel heißt nicht *ranken*, sondern **zitiert werden**: eine der Quellen zu sein, die die KI liest, ihr vertraut und namentlich nennt. Dafür stehen zwei Disziplinen: **GEO** (Generative Engine Optimisation) und **AEO** (Answer Engine Optimisation). Antwortmaschinen bevorzugen Inhalte, die klar strukturiert, faktisch fundiert, maschinenlesbar und nachweislich vertrauenswürdig sind. Konkret heißt das: saubere strukturierte Daten (JSON-LD), eine `llms.txt`-Datei, starke **E-E-A-T**-Signale, eindeutige **Entitäten** und antwortförmig aufbereiteter Content. SEOWing operationalisiert all das per MCP-Tool — Schema, `llms.txt`, E-E-A-T und Entitäten — statt mühsamer Handarbeit. Am schwersten ist die **Messung**: KI-Antworten sind nicht-deterministisch und ohne klassischen Rankingreport. Der pragmatische Weg: kontrollierbare Leitindikatoren instrumentieren und diszipliniertes Prompt-basiertes Zitat-Testing auf festen Fragensets fahren. Mehr unter **seowing.io**.