

AI-Powered SEO (AEO & GEO) — Workbook

This workbook turns the course into action on your own site. Work through one section per module: map how engines source your niche, build a real question inventory and answer-ready outlines, mark up and trust-proof your pages, then stand up a monthly measurement loop. Fill the worksheets and templates with your own data, and leave every total, score, and calculated cell blank until you work it out yourself.

How AI Answer Engines Actually Choose Sources

Learn how the four major engines cite sources in your niche so you can sequence your effort for the fastest wins.

Exercise: Run the four-engine citation test

Pick one high-value buying or how-to question in your niche and ask it, identically, in Google (with AI Overviews), ChatGPT search, Perplexity, and Gemini. Record who gets cited so you can see where a brand like yours has the best near-term odds. Do not edit the answers; copy what you see.

- What exact question did you ask, and what is the buyer intent behind it?

- For each engine, which 3 to 8 sources were cited, and were you or a competitor among them?

- Which engine surfaced the freshest or most specialist sources, and which leaned on big established brands?

- Based on the pattern, which engine will you target first, and why?

Worksheet: RAG-lever opportunity map

For one priority page, judge where you are weakest in the retrieve-select-cite pipeline so you fix the binding constraint first. Complete one row per page you care about.

Page URL

Target question / prompt

Indexed in Google? (Y/N)

Indexed in Bing? (Y/N)

Covers the fanned-out sub-questions? (Y/N)

Has a clean extractable answer passage? (Y/N)

Has trust signals (author, sources, schema)? (Y/N)

Weakest lever (retrieve / select / cite)

Checklist: Engine-readiness foundations

- [] The page is indexed in BOTH Google and Bing (checked, not assumed)
- [] robots.txt does not block GPTBot, PerplexityBot, or Google-Extended unless intended
- [] An accurate Organization entity exists with sameAs links
- [] You have identified which engine offers your fastest near-term wins
- [] You can name which RAG step (retrieve / select / cite) each planned tactic improves

Prompt-Intent Mapping and Answer-Ready Content

Build a real inventory of the conversational questions your audience asks AI, then shape each answer to its intent.

Exercise: Build a question inventory with seed-and-fan

Take one seed topic tied to a page you own, generate the real prompts buyers ask, and cluster them. Use AlsoAsked or AnswerThePublic plus an AI engine, then sanity-check demand in a keyword tool.

- What is your seed topic and the single page it maps to?

- List 20 to 30 conversational questions (who/what/why/how/best/vs) real buyers would ask about it

- Tag each question by intent: informational, commercial, comparison, or transactional

- Which constraints recur (budget, team size, integration, location), and which question clusters can one page answer together?

Worksheet: Intent-to-format mapping

Decide the answer shape each question demands before you write, so engines can extract it. Fill one row per question from your inventory.

Question (phrased as a person would ask it)

Intent type (definition / how-to / comparison / best-of / yes-no)

Required answer format (short definition / numbered steps / table / shortlist / yes-no + reason)

Heading text to use on the page

Target page URL

Exercise: Rewrite a section for extractability

Take one weak section of an existing page and rewrite it using the inverted pyramid and chunking. Keep both versions so you can compare.

- Paste the original section, then rewrite it so the first sentence states the full answer

- Add one specific statistic with an inline source, and one credible quote or reference if the claim is contestable

- Replace pronouns that point to earlier text with the named subject so the chunk stands alone

-
- Is the rewritten section shorter per point and denser in facts than the original? Note what changed.
-

Checklist: Answer-ready content standards

- Each user question appears as an H2 or H3, phrased the way people ask it
- The first sentence under each heading is the direct answer
- Paragraphs are 2 to 4 sentences; lists and tables are used for any set of items
- Concrete numbers, dates, and named entities replace vague claims
- Each passage makes sense if lifted out of the page (no orphan pronouns)
- At least one cited statistic or credible source supports contestable claims

Structured Data and Entities for LLM Citation

Mark up pages with the right schema and build the entity and trust signals that make engines confident enough to cite you.

Worksheet: Schema plan per page

Choose the structured-data types each priority page needs and confirm they match visible content. Complete one row per page, then validate before publishing.

Page URL

Primary content type (article / FAQ / product / how-to)

Schema types to add (FAQPage / Article / Organization / Product / HowTo / BreadcrumbList)

Author Person schema present? (Y/N)

sameAs links included? (Y/N)

Validated in Rich Results Test? (Y/N)

Matches on-page content exactly? (Y/N)

Exercise: Map your brand as an entity

Make your brand an unambiguous, well-connected entity so engines can reason about it. Audit your current entity footprint and plan the gaps.

- Is your name, address, and phone (NAP) identical across your site and major directories? List any mismatches
 - Which entity-defining sources mention you (Wikipedia, Wikidata, Crunchbase, industry directories), and which are missing?
 - What sameAs links does your Organization schema include, and which credible profiles should you add?
 - Which other recognized entities in your niche do you currently appear alongside, and how could you earn more co-mentions?
-

Worksheet: Pillar-and-cluster topical map

Plan one pillar page and its cluster pages so you cover a topic's full question space and feed query fan-out. List the pillar, then one row per cluster page.

Pillar page topic and URL

Cluster page sub-question

Cluster page URL

Intent type

Internal link to pillar in place? (Y/N)

Last updated date

Checklist: E-E-A-T and trust pre-flight

- A real, named author with a credentialed bio and Person schema is attached
- At least one piece of original, first-hand evidence (photo, test result, data) is on the page
- Every factual or statistical claim cites a primary, dated source
- A visible last-updated date is present and stale facts are refreshed
- About, Contact, and editorial/review policy pages are easy to find
- For products/services, reviews or third-party reputation signals are shown

Measuring, Tracking, and Improving AI Visibility

Stand up an honest measurement baseline and run a monthly loop that compounds citations over time.

Worksheet: Priority prompt set

Define a stable set of 15 to 25 prompts that matter to your business so your share-of-voice trend is comparable month to month. Complete one row per prompt; keep this list fixed once set.

Prompt text (as a buyer would type it)

Intent type

Mapped page URL

Primary competitor to watch

Engines to check (AI Overviews / ChatGPT / Perplexity / Gemini)

Exercise: Audit a page for AI-citation readiness

Run the six-dimension audit on one high-intent page, scoring each area from 0 (absent) to 3 (excellent). Fix the lowest-scoring dimension first. Do the totals yourself; do not skip the scoring.

- Score retrieval, intent match, extractability, structured data, trust, and entity clarity from 0 to 3 each
-

- Add the six scores for a total out of 18, and judge: 14+ is citation-ready, below 10 needs structural work
-
- Which single dimension scored lowest, and what is the specific fix?
-
- What is the page's before score, and what do you expect after the fix?
-

Checklist: Monthly AEO and GEO loop

- Ran the priority prompt set and updated the share-of-voice tracker
- Listed prompts where a competitor is cited and you are not, with their cited page noted
- Audited the competing page on the six-dimension rubric
- Fixed the weakest dimension and published the change
- Requested indexing in Google Search Console and Bing Webmaster Tools
- Confirmed no unwanted AI crawler is blocked in robots.txt
- Logged before-and-after audit scores to track what works

Your Action Plan

1. Run the four-engine citation test on one priority question and decide which engine to target first
2. Confirm your priority pages are indexed in both Google and Bing and that GPTBot, PerplexityBot, and Google-Extended are not blocked
3. Build a question inventory for one topic using seed-and-fan, then cluster the questions into pages
4. Map each question to an intent and a required answer format, and turn the questions into on-page headings
5. Rewrite your three highest-intent pages with inverted-pyramid openings, chunked passages, and at least one cited statistic each
6. Add and validate the right schema (FAQPage, Article, Organization, Product as relevant) on those pages
7. Fix your brand entity: consistent NAP, an accurate Organization schema with sameAs links, and presence on entity-defining sources
8. Attach a credentialed author, original first-hand evidence, and dated source citations to each page for E-E-A-T
9. Define a stable set of 15 to 25 priority prompts and record a baseline share-of-voice across the four engines
10. Audit each priority page on the six-dimension rubric, fix the lowest score first, then run the measure-diagnose-audit-fix loop monthly

