GEO White Paper for Export Manufacturers

Helping AI search find your factory correctly

For owners and the export and marketing teams of manufacturers selling abroad: what GEO is, how AI answers choose sources, how to measure in a way others can check, and a 90-day roadmap.

Version	v1.0 (2026-11-27)
Audience	Owners, export and marketing teams of manufacturers selling abroad
Nature	A synthesis of public sources and a method framework. No client cases or names; no promise of rankings, AI recommendations or inquiry volume
Data policy	External figures cite their source and access date; inferences, estimates and worked examples are labelled

Summary

Overseas buyers increasingly ask an AI assistant first and decide which website to open afterwards. GEO (Generative Engine Optimization) addresses one concrete question: **when a buyer asks where to find a reliable maker of some product, does the AI's answer mention your factory, describe it correctly, and send the buyer to your website?**

- **GEO is not buying rankings.** Google's own guidance says that optimizing for its generative AI features is still SEO, and that tactics such as chunking content, llms.txt files and inauthentic mentions can be ignored [2].
- **AI answers vary; a single screenshot is not a measurement.** A study that repeated the same prompts found that the same brand list almost never reappears [4]. Track the appearance rate over a fixed question bank, with a confidence interval.
- **Source preferences differ by platform.** The most-cited sources are not the same across platforms [6]. Third-party web mentions correlate more strongly with brand visibility than backlink counts do [5] (correlation, not causation).
- **Measurement must be checkable.** Fixed question bank, fixed conditions, enough samples, Wilson intervals, an untouched control group, and the raw answers kept on file.
- **Roll-out takes four stages over roughly 90 days.** Baseline, fixes, content and corroboration, re-test. Each stage has deliverables and acceptance checks; product facts are approved by the company itself.
- **Compliance is a boundary, not an appendix.** No fake reviews, no paid posts, no fabricated mentions; structured data matches visible content; AI-generated content is reviewed by a person.

> **How to read this**
>
> For the concept, read sections 1 and 2. For acceptance or contracts, read 3 and 5. To get started, read 4 and prepare the checklist in section 6 first.

Contents

01 1 What GEO is, and how it relates to SEO
02 2 How AI answers choose sources
03 3 How to measure: fixed question bank, Wilson intervals, control group
04 4 A 90-day roadmap
05 5 Compliance boundaries
06 6 Before you start: eight items on the company side
07 Sources
08 Disclaimer

1 What GEO is, and how it relates to SEO

The term comes from a 2023 paper by Aggarwal et al. (KDD 2024). It defines a generative engine as one that retrieves several sources and has a language model synthesize an answer, and it reports on its own benchmark that certain rewrites of content can raise visibility in the answer by up to about 40% [1]. That figure was measured in a controlled setting and describes share within a single answer. It is not traffic, clicks or orders, and it should not be read as an expected result for a company website.

Aspect	SEO	GEO
Goal	Position and clicks in search results	Being mentioned, cited and described correctly in AI answers
Unit	Keywords and pages	The buyer's question (intent) and the sources behind the answer
Stability	Rankings are relatively stable and can be tracked daily	Answers vary at random; repeated sampling is needed
Main levers	Crawlability, relevant content, links	Accurate, crawlable facts on your site, consistent facts everywhere, verifiable third-party corroboration
Acceptance	Rank, impressions, clicks	Appearance rate, citation rate and factual accuracy on a fixed question bank, with intervals

They are not substitutes. Google's guide states that from its perspective, optimizing for generative AI search is optimizing for the search experience, and so still SEO [2]. In practice the foundation (crawlable, indexable, truthful and useful content) overlaps with SEO. What GEO adds is measurement at the answer layer and management of third-party mentions and entity consistency.

What this means for an export manufacturer

Buyers' questions usually carry conditions: product category, material, certification, minimum order, origin, lead time. An AI answer often lists only a handful of suppliers, and a factory whose information is missing or inconsistent is more likely to be left out or described wrongly. What a factory can do is publish checkable facts on its website and keep the same facts consistent in directories, certification-body pages and trade-show materials.

2 How AI answers choose sources

2.1 Retrieve, synthesize, and expand the question

Google explains that its AI features retrieve relevant pages from the search index, generate a response from them and show links. Query fan-out means the model generates a set of concurrent, related queries to fetch more results [2]. One buyer sentence can therefore become several sub-queries, so your pages need to answer not only the original wording but its facets: specifications, certification, compatibility, commercial terms.

2.2 First make sure crawlers are not blocked

OpenAI documents OAI-SearchBot and GPTBot as independent robots.txt settings: a site can allow the first to appear in search results while disallowing the second for training [3]. ChatGPT-User is a fetch triggered by a user

during a conversation. Checking whether robots.txt, the CDN or a firewall blocks search crawlers is the cheapest first step.

2.3 Platforms prefer different sources

Platform (Profound sample, 2024-08 to 2025-06)	Top three sources by share of all citations
ChatGPT	Wikipedia 7.8%, Reddit 1.8%, Forbes 1.1%
Google AI Overviews	Reddit 2.2%, YouTube 1.9%, Quora 1.5%
Perplexity	Reddit 6.6%, YouTube 2.0%, Gartner 1.0%

Source: Profound public article, covering about 680 million citations it tracked [6]. The sample reflects that company's monitoring scope and ends in mid-2025; platform behaviour changes, so treat it as direction, not current state.

Another often-quoted observation is from Ahrefs. Across 75,000 brands, branded web mentions had a Spearman correlation of 0.664 with brand visibility in AI Overviews, branded anchors 0.527, and backlink count only 0.218 [5]. Ahrefs itself stresses that this is correlation. What can be borrowed is the direction: verifiable third-party mentions deserve more attention than raw link counts.

> **Inference (not a research finding)**
>
> A prudent approach for a factory is to work on three kinds of source together: primary facts on its own site (specifications, certifications, capacity, contact details), industry directories and certification-body pages, and verifiable trade-show and media coverage. Do not bet visibility on a single source type, because platform preferences change.

3 How to measure: fixed question bank, Wilson intervals, control group

3.1 Why one screenshot is not a measurement

SparkToro and Gumshoe asked AI tools the same prompts about 3,000 times and concluded that the same brand list reappears with a probability below one in a hundred, and the same order even less often [4]. A rank position therefore carries almost no meaning. What can be compared is the share of questions, in a defined set, where you appear.

3.2 Question bank: traceable sources, frozen versions

- **Traceable sources:** on-site search terms, Search Console queries rewritten as conversational questions, sales inquiries and support logs, trade-show questions, lists the client has confirmed in writing. The China Advertising Association standard T/CAAAD 046-2026 likewise requires samples to have an explainable, traceable source [7].
- **Layering:** report branded questions (to check facts) separately from non-branded questions (to observe recommendations), and group intents into knowledge, recommendation, comparison and verification.
- **Phrasing variants:** write three to five natural phrasings per intent, and a set for each market language.
- **Frozen versions:** the core question set used for acceptance is confirmed by both parties at kick-off and not changed once acceptance starts; any change needs written confirmation and a note of its effect on results [7].

3.3 Metric definitions

Metric	Definition	Note
Mention rate	Valid answers that mention the brand ÷ all valid answers	Tools use different denominators; do not compare across tools
Recommendation rate	Answers placing the brand in a pre-defined recommended set ÷ valid answers to recommendation or comparison questions	Define the set up front, for example the top three
Citation rate	Valid answers containing at least one target URL ÷ all valid answers	Separate pages fetched from pages shown as citations
Factual accuracy	Correct fact items ÷ fact items checked	Check against the fact sheet the company approved; log each wrong sentence
Average position	Diagnostic only	Position is very unstable; not recommended as a KPI

Failed requests, degraded answers and refusals are counted separately, not silently treated as “not mentioned”.

3.4 Sample size and Wilson intervals

An appearance rate is a binomial proportion and should carry a confidence interval. The Wilson interval behaves better than the usual normal approximation with small samples and with proportions near 0% or 100% [10]. In the worst case (proportion 50%, simple random sampling), 95% confidence with a margin of ±10 percentage points needs about 97 answers and ±5 points needs about 385 (computation: $n \geq 1.96^2 \times 0.25 \div \text{margin}^2$). AI answers are not fully independent random draws, so the real requirement is usually higher.

Worked example (assumed numbers, to show the calculation)	Observed rate	95% Wilson interval
2 appearances in 20 answers	10.0%	about 2.8% to 30.1%
20 appearances in 100 answers	20.0%	about 13.3% to 28.9%
60 appearances in 300 answers	20.0%	about 15.9% to 24.9%

The same 20% narrows from about 16 points wide to about 9 points wide when the sample grows from 100 to 300. Two numbers whose intervals overlap cannot be ranked against each other.

3.5 Control group: separating your change from platform change

- **Keep an untouched control:** about a fifth of the questions, or a set of similar pages, stay unchanged during the observation period. The gap between changed and control groups is a closer estimate of the effect of your change.
- **Record conditions:** platform and model version, date and time, request text, test account state, location, web access and conversation context, repeat count, raw answer, cited links and screenshots. This is also the minimum record that T/CAAAD 046-2026 asks for in a baseline assessment [7].
- **Measure the surface buyers actually see:** a chat interface and an API can behave differently, so agree the measurement surface in advance.

- **Do not hide anomalies:** the standard says not to selectively conceal anomalous samples, failed samples or changes of definition, and not to use alarmist or exaggerated claims to close a deal [7].

3.6 Traffic and pipeline layer

In GA4 you can build an AI-assistant channel in a custom channel group and order it above Referral, otherwise those sessions are assigned to Referral first [11]. Server logs can separate training, search-indexing and user-triggered crawlers [3]. Adding a “How did you find us?” field to the inquiry form recovers traffic that carries no referrer. These are observed sources, not proof of causation.

4 A 90-day roadmap

The following is a reference path that starts with one market and one product category. Timings are estimates; they depend on how quickly product facts can be approved and on access to the website back end.

Stage	What happens	Deliverables	Acceptance
Weeks 1-2 Baseline	Confirm market and category; build and freeze the question bank; first sampling on chosen engines; check crawler access, sitemap, titles and redirects	Question bank v1, baseline report, issue list	Question bank confirmed in writing; baseline shows sample size and intervals
Weeks 3-6 Fixes	Fix missing and inconsistent facts on the site; complete specifications, certifications and contact details; Organization and Product structured data match visible content; align directory profiles	Change list and release log	Spot-check each item; keep original values before release
Weeks 7-10 Content and corroboration	Write buyer Q&A; and specification pages for gaps in the question bank; make information consistent on verifiable industry directories, certification bodies and trade-show materials; every item approved by the company	Content and corroboration list	Product facts approved by the company; external placements logged with link and whether paid
Weeks 11-13 Re-test	Re-run the frozen question bank; compare with the control group; summarize inquiry sources; decide the next round	Re-test report, next-round proposal	Intervals shown beside the control group, including anomalous samples

No improvement is promised for any week. If the intervals overlap at re-test, the finding is written as “no distinguishable change”, not as whichever number looks best.

5 Compliance boundaries

- **No fabrication:** no fake reviews or ratings, no paid posts, no inauthentic mentions. Google’s AI optimization guide lists inauthentic mentions among tactics that can be ignored and should not be pursued [2], and buying or selling links to manipulate ranking is named in its spam policies [8].
- **Structured data matches the page:** marked-up content must be real information readers can see [9].
- **AI-generated content is reviewed by a person:** generating many low-value pages to manipulate ranking is “scaled content abuse” [8]. Product facts, certifications and numbers come only from material the company has approved.

- **Honest reporting:** state the method, sample size, intervals, failed samples and changes of definition [7].
- **Data minimization and cross-border handling:** collect only what the service needs, tell people in advance how data is handled and transferred, and follow applicable personal-data rules. The written contract and legal advice govern the details.
- **No guaranteed outcomes:** AI and search platforms change their answers and policies, so no service should promise rankings, AI recommendations or inquiry volume.

6 Before you start: eight items on the company side

1. The target market, language and first product category, and a technical or sales contact who can approve product facts.
2. Edit access to the website back end, or contact details of the web vendor who will make changes.
3. A list of certifications and qualifications (certificate number, validity, issuing body) and which may be made public.
4. Genuine inquiries and common buyer questions from the last 6 to 12 months (identities removed).
5. Public wording for minimum order, lead time, supply regions and payment terms, identical across the site.
6. URLs of existing directory, trade-show and platform profile pages, to check wording consistency.
7. Who answers inquiries that arrive through AI channels, and the response time limit.
8. Written acceptance terms: question-bank version, engines and measurement surface, repeat count, observation period.

Sources

Accessed 2026-11-27. “Opened and checked” means the original was opened and the cited content checked while writing; “secondary” means the original could not be opened and the figure relies on public reporting, so readers should verify it.

[1] Aggarwal et al., GEO: Generative Engine Optimization, arXiv:2311.09735 (KDD 2024). Abstract opened and checked.
https://arxiv.org/abs/2311.09735
Used in section 1 for the definition and the scope of the “up to about 40%” figure.

[2] Google Search Central, Optimizing for generative AI features on Google Search (last updated 2026-07-10). Opened and checked.
https://developers.google.com/search/docs/fundamentals/ai-optimization-guide
Used in sections 1, 2 and 5: still SEO, query fan-out, llms.txt and inauthentic mentions can be ignored.

[3] OpenAI, Overview of OpenAI Crawlers. Opened and checked.
https://developers.openai.com/api/docs/bots
Used in sections 2 and 3: OAI-SearchBot, GPTBot and ChatGPT-User.

[4] SparkToro, study on how inconsistently AI tools recommend brands (secondary). The site returned 403 to automated access, so the original was not opened; the figure follows Search Engine Land’s report.
https://sparktoro.com/blog/new-research-ais-are-highly-inconsistent-when-recommending-brands-or-products-marketers-should-take-care-when-tracking-ai-visibility/
Used in section 3.1; defer to the original.

[5] Ahrefs, correlation study of brand visibility in AI Overviews (75,000 brands). Opened and checked.
https://ahrefs.com/blog/ai-overview-brand-correlation/
Used in section 2.3; coefficients are correlations, not causation.

[6] Profound, AI platform citation patterns (about 680 million citations, 2024-08 to 2025-06). Opened and checked.
https://www.tryprofound.com/blog/ai-platform-citation-patterns
Used for the table in section 2.3; the sample is that company’s monitoring scope.

[7] China Advertising Association, T/CAAAD 046-2026 Requirements for Generative Engine Optimization Services (issued 2026-10-13). Clauses 7.2.1, 7.2.3, 7.4.2 to 7.4.4 and 12.1.2 opened and checked.
https://digitalad.china-caa.org/TechLab/pdf/TCAAAD%20046-2026.pdf
Used in sections 3 and 5; a group standard, not binding regulation.

[8] Google Search Central, Spam policies for Google web search. Opened and checked.
https://developers.google.com/search/docs/essentials/spam-policies
Used in section 5: link buying and selling, scaled content abuse.

[9] Google Search Central, General structured data guidelines.
https://developers.google.com/search/docs/appearance/structured-data/sd-policies
Used in section 5: markup must match visible content.

[10] Wilson score interval (statistical method).
https://en.wikipedia.org/wiki/Wilson_score_interval
Used in section 3.4; sample sizes and intervals are this paper’s own calculations, not survey data.

[11] Google Analytics Help, Custom channel groups.
https://support.google.com/analytics/answer/13051316
Used in section 3.6; channels are assigned in order and the first match wins.

Disclaimer

This white paper was written by wantone Search Consulting (operated by Taiwan Mobile Net, Taiwan unified business number 38469314). It is general information only and is not legal, tax, investment or other professional advice.

It promises no search ranking, no AI assistant recommendation or citation, and no inquiry volume or revenue. The answers, source preferences and policies of AI and search platforms, including Google, OpenAI and Perplexity, change with time, model and phrasing, and are outside our control.

External figures come from the public sources listed, which differ in definition, sample and date, and some are correlational. Items marked as inference, estimate or worked example are numbers derived or assumed in this paper, not survey results. Secondary sources have not each been re-verified by us.

External links lead to third-party sites whose content and availability we do not control. The standards, platform guidance and statistics cited may have been updated; check the current version before relying on them.

This paper contains no client cases or client names. Copyright belongs to wantone and the original authors; please cite the source. Contact: contact@search.com.tw.