AI Search Jobs 2026: what 867 postings ask for, and what the evidence shows

Fernando Angulo

Original analysis: October 12, 2026. Corpus cutoff: September 22, 2026, freeze 1.0. PDF edition 1.0: November 4, 2026.

Personal preliminary research. Not peer-reviewed. This PDF preserves the original October 2026 analysis and source-availability context; its formatting and bibliography do not represent a new review of the data. The study uses public job-posting sources, not proprietary Semrush or Adobe datasets.

Abstract

This preliminary descriptive study examines a convenience sample of 1,003 unique job postings frozen on September 22, 2026: 867 classified as current at the cutoff and 136 historical. It distinguishes titles, work dimensions, evidence states, and the scope of commercial expectations. Results describe the wording and coding of recovered postings, not labor-market prevalence, hires, or demonstrated commercial return. Review was AI-assisted; 944 records retain pending final-adjudication statuses.

Suggested citation

Angulo, F. (2026). *AI Search Jobs 2026: what 867 postings ask for, and what the evidence shows*. Preliminary analysis, PDF edition 1.0. https://fernandoangulo.com/leadership/ai-search-jobs

Analysis

Personal research by Fernando Angulo. Preliminary analysis of the corpus frozen on September 22, 2026. The main sample contains 867 postings classified as current at that cutoff, from 701 employers. A separate historical cohort contains 136 postings. Current describes the cutoff classification; it does not guarantee a vacancy remains open today.

AI search work appears beyond job titles carrying its labels. **495 of the 867 current postings (57.1%) contain no explicit AI search label in their title**, despite meeting the functional inclusion criteria. The remaining 372 (42.9%) match the study's title dictionary.

My reading is that searching only for a GEO Manager misses some of the work employers in this sample describe. GEO means *Generative Engine Optimization*; AEO means *Answer Engine Optimization*. The function can also sit within SEO, content, growth, or marketing leadership. This finding describes the corpus, not the size of the labor market.

The second signal is the distance between visibility and business impact. Presence is coded in 683 current postings and commercial impact in 332. However, only two of those 332 have a commercial-impact code labeled with the exact AI search scope. Much of the commercial language belongs to the wider role. Assigning revenue expectations and demonstrating revenue attributable to AI search are different observations.

The sample: from 1,466 records to 867 postings in the main analysis

The unit is a unique posting or requisition, not a URL, location, hire, or employed person. A multi-location posting counts once. Deduplication prioritizes employer and requisition identifier, supported by canonical URLs and checks of title, location, and content.

Frozen registry stage	Count	Use
Physical records collected	1,466	Includes aliases and candidates not admitted
Source aliases removed	26	Another representation of the same unit
Candidate units after removing aliases	1,440	Discovery universe
Core postings admitted	1,003	867 current and 136 historical
Main cohort: official current source at cutoff	867	Main denominator
Historical cohort: official historical evidence	136	Separate denominator
Candidates awaiting resolution	344	Outside the results
Exploratory adjacent records	85	Outside the core
Excluded records	8	Outside the core

Discovery was targeted and multilingual, combining GEO, AEO, AI Search, and AI Visibility terms with conventional titles. Sources included employer pages and recruitment systems such as Greenhouse, Lever, Ashby, Workable, and SmartRecruiters. Directories and intermediaries helped discover candidates; they did not replace primary evidence for inclusion.

A posting had to describe substantive work on commercial generative search: visibility, inclusion, citations, recommendations, or activities enabling them. Generic AI references, model-engineering positions without this function, and geospatial uses of GEO were insufficient. The protocol distinguishes official current sources, official historical sources, and discovery sources.

This is a **convenience sample**, shaped by queries, languages, access, and source recovery. It is neither a probability sample nor a census. All eight files match their frozen SHA-256 hashes. I recalculated the tables from the frozen registry. The sample register, cohort tables, and methodology make the calculation reviewable.

Titles: 57.1% fall outside a search based only on explicit labels

Title classification	Current, n = 867	Historical, n = 136
Explicit term from the AI search dictionary	372 (42.9%)	58 (42.6%)
No explicit dictionary term	495 (57.1%)	78 (57.4%)

This is a lexical classification. It recognizes GEO, AEO, AIO, generative engine, answer engine, AI Search, AI Visibility, and LLM Visibility according to the published regular expression. It does not measure specialization or seniority. No explicit term means that the title does not match this dictionary, not that the posting has no AI work.

The proportions are similar across cohorts, but these cohorts are not comparable time windows. The table describes two groups of recovered postings; it does not establish a change in title usage.

For a marketing leader, the practical implication is to inventory tasks before titles. An employer can assign AI search work without creating a dedicated position. The sample documents that coexistence, without establishing the dominant organizational structure across the market.

Nine dimensions of work and expectations

Each posting can receive several codes, so percentages do not add to 100%. The table counts each code once per posting, across any role scope and any of four evidence states: responsibility, target, metric, or required prior result. **Not every code is a task, and not every code is specific to AI search.**

Coded dimension	Current, n = 867	Historical, n = 136
O0. Operating capability: processes, teams, governance	679 (78.3%)	76 (55.9%)
O1. Eligibility: crawling, retrieval, interpretation	465 (53.6%)	67 (49.3%)
O2. Presence or mention	683 (78.8%)	113 (83.1%)
O3. Citation or explicit attribution	318 (36.7%)	60 (44.1%)
O4. Selection or recommendation	84 (9.7%)	17 (12.5%)
O5. Qualified visit	147 (17.0%)	24 (17.6%)
O6. Initial conversion	295 (34.0%)	45 (33.1%)
O7. Commercial impact	332 (38.3%)	41 (30.1%)
O8. Accuracy, trust, and risk	145 (16.7%)	34 (25.0%)

Eligibility includes making content crawlable, retrievable, or interpretable. Presence means the brand appears; citation requires explicit attribution. Recommendation involves selection or comparison. Initial conversion can be a signup, inquiry, or purchase. Commercial impact covers pipeline, revenue, retention, or economic efficiency.

Restricting the count to the **responsibility** evidence state leaves 679 current postings with operating capability and 453 with eligibility. Presence falls to 22 and citation to eight because these codes appear mainly as metrics or targets. Describing all 683 presence-coded postings as execution responsibilities would erase that distinction.

I also checked employer concentration. Counting an employer once per dimension, presence appears in 571 of 701 employers (81.5%), operating capability in 559 (79.7%), eligibility in 396 (56.5%), and citations in 289 (41.2%). This sensitivity check prevents employers with multiple requisitions from explaining the pattern alone. The denominator still consists of employers found through targeted discovery.

Scope changes the interpretation of commercial expectations

The next table applies the exact scope filter `ai_search`. It excludes general-role codes, combined organic-and-AI-search labels, and other unnormalized scopes. This is a conservative filter, not a complete measure of every AI-related reference.

Dimension	Any role scope	Exact AI search scope	Denominator
Presence	683 (78.8%)	681 (78.5%)	867 current postings
Citation	318 (36.7%)	318 (36.7%)	867 current postings
Qualified visit	147 (17.0%)	36 (4.2%)	867 current postings
Initial conversion	295 (34.0%)	15 (1.7%)	867 current postings
Commercial impact	332 (38.3%)	2 (0.2%)	867 current postings

My interpretation is that the sample documents visibility more clearly than AI-specific commercial attribution. That reading combines employer language with coding decisions. It does not show that employers lack attribution or that visibility cannot contribute to sales. An absent code in a posting does not establish an absent practice.

Across the 1,003 core postings, **no record contains an achieved commercial result** in the field dedicated to that evidence. The corpus contains expectations, metrics, and experience requirements, not observed hiring or business outcomes. It studies how work is defined; it does not demonstrate its return.

Four postings show how the work fits together

These cases illustrate cutoff coding, without representing their sectors. Links point to the original sources. Some no longer return readable text; the cutoff evidence remains referenced in the registry.

Stripe, Head of AEO & SEO. The posting combines team leadership, architecture, content, and measurement. It distinguishes visibility, citations, signups, and economic impact from organic channels. The mandate requires coordination across product, engineering, and marketing; the posting provides no achieved results. The Stripe source remained readable during this review.

Aiven, Technical AEO Manager. The role combines retrievable content structure, a prompt panel, citations, and connections to commercial data. It also asks for attribution caveats. It illustrates why a technical function needs access to documentation, engineering, and sales. The Aiven posting remained readable during the review.

Virgin Voyages, Sr. Discoverability (GEO/SEO) Manager. Frozen coding records cross-functional strategy, technical readiness, presence, citations, and recommendation for travel decisions. Business objectives carry a general-role scope. The original source returned no verifiable text during this review; this case comes from the frozen registry.

Smarsh, Generative AI & SEO Strategist. The registry records structured data, entities, visibility, and citations alongside MQLs and pipeline for the wider organic role. The useful distinction is between AI search diagnosis and overall commercial contribution. The original source could not be recovered during this review; I do not claim the vacancy remains open.

Review coverage and the remaining adjudication gap

The review register lists all 1,003 core postings as completing materiality review and second outcome coding. Coder identifiers document AI-assisted work. This is not independent human expert peer review, and no inter-coder agreement measure such as kappa is recorded.

One discrepancy remains visible: **944 records retain a second-coding-complete, final-adjudication-pending status**; 59 have completed-adjudication statuses. Of the pending records, 844 are current and 100 historical. Reproducing the totals validates the arithmetic; it does not resolve the validity of every label.

These tables are therefore preliminary results from the frozen registry. A final report requires reconciliation of these statuses with final decisions and review of the affected coding evidence. Skills, compensation, and seniority frequencies also require sufficient normalization and coverage before their denominators can support publication.

What a CMO can decide from this evidence

The corpus helps define a mandate. My recommendation is to separate four responsibilities: maintain retrievable and accurate evidence, observe presence and citations, coordinate interventions, and evaluate commercial connections with their attribution limits.

To prepare the mandate, record five fields: buyer decision, evidence to maintain, specialists and access required, signals to observe, and the result criterion. Compare them with existing team capabilities before choosing to hire, train, or reassign.

A job description should state who can act on documentation, content, engineering, and data. If revenue is an expectation, provide the instrumentation and explain what will count as evidence of a result. A visibility screenshot does not settle that requirement.

The topic ownership map helps identify the questions and evidence a brand should cover. The reporting model separates observation, interpretation, and business outcomes. Together they help a team decide whether to hire, train, or redistribute work. These 867 postings provide design references; they do not automatically justify another vacancy.

Frequently asked questions

Does the sample contain 1,003 open vacancies?

No. It contains 867 classified as current at the September 22, 2026 cutoff and 136 historical postings. It is not a count of 1,003 vacancies open today or 1,003 hires.

How many titles carry an explicit AI search label?

372 of 867 current postings, or 42.9%, match the published dictionary. The other 495 meet the functional inclusion criteria without matching that dictionary in their title.

Which dimensions are most frequent?

Presence, 683 of 867 (78.8%), and operating capability, 679 (78.3%). Codes combine responsibilities, targets, metrics, and required experience, so their evidence state matters.

Do these postings demonstrate commercial return from AI search?

No. They document expectations. The achieved-commercial-results field is empty across all 1,003 core postings.

Does counting employers change the pattern?

Presence appears in 571 of 701 employers (81.5%), operating capability in 559 (79.7%), eligibility in 396 (56.5%), and citations in 289 (41.2%). The pattern is not explained solely by employers with several requisitions.

Is this a final study reviewed by human experts?

No. It is a reproducible preliminary analysis with AI-assisted review and pending adjudication statuses that must be reconciled before finalizing the report.

Before approving the next vacancy, complete the responsibility, access, and results-evidence map with your team. That conversation will make the title you choose more useful.

Methodology

Version: editorial reanalysis, October 11, 2026. Source cutoff: September 22, 2026, freeze 1.0. Personal research, public job-posting sources, no proprietary employer datasets. Preliminary results.

Design and eligibility

Descriptive job-content study using targeted multilingual discovery, not a probability sample, census, blinded preregistration, or representative labor-market estimate. The protocol followed an exploratory pilot. Searches combined explicit AI search terminology and conventional titles; employer career pages and ATS sources were primary evidence. Intermediaries and directories were discovery leads. Eligibility required material commercial generative-search responsibilities; generic AI work and geospatial GEO were insufficient.

Unit: unique requisition. Employer plus stable posting/requisition identifier is preferred for deduplication; canonical URL and manual company/title/location/body comparison support it. Multiple locations and source aliases do not constitute additional jobs or hires.

Flow: 1,466 physical records, minus 26 aliases = 1,440 candidate units. Core = 1,003 (867 current-tier-A, 136 historical-tier-B); remaining units = 344 pending, 85 exploratory-adjacent, eight excluded. These categories reconcile exactly. Current is the source-status classification at the freeze, not confirmation of present-day availability. Historic recovery is not a comparable earlier sampling window.

Calculation

The recalculation uses the frozen registry and asserts all eight source hashes in its manifest. The editorial review retains the reproduction script and source hashes. `cohort-analysis.json` includes the hashes and title expression. `sample-register.csv` lists IDs, cohorts, employers, titles, source URLs, source tier, capture date where recorded, posting date where recorded, and coding status. Raw job-description text is not redistributed.

Published title classifier reproduces the reviewed snapshot, not the older exploratory script:

```
(?:\bgeo\b|\baeo\b|\baio\b|generative engine|answer engine|ai[- ](?:search|visibility)|llm[- ]
visibility)
```

Case-insensitive lexical matching, not a validated specialist-title classifier. Explicit/current: 372; other/current: 495. Explicit/historical: 58; other/historical: 78. Pooled title total 430 is a reconciliation check only; the article reports cohort denominators separately.

Outcome counts use the set of `outcome_code` values within `outcomes_legacy` per row. Each code contributes at most once per requisition. Structured codes O0-O8 represent operating capability, eligibility, presence/mention, citation, recommendation, qualified visit, initial conversion, commercial impact, and accuracy/risk. They are nonexclusive. Counts combine responsibility, target, metric, and required-prior-result states unless filtered explicitly. A code is not necessarily a responsibility, an AI-specific outcome, or an achieved result.

The exact `scope == ai_search` filter excludes combined organic-and-AI labels and numerous heterogeneous, unnormalized scope labels. It is conservative, not an exhaustive AI-related estimate. This is particularly important for O5-O7. Broader commercial duties cannot be assigned to AI search alone. Zero achieved-commercial-result fields mean no such evidence is documented in this registry, not that no company achieved results.

Employer sensitivity counts each employer label at most once per code among the 701 current employers. It uses the frozen employer identities rather than a new parent-company consolidation. Percentages equal count/cohort denominator, rounded to one decimal; rounding uses Python's standard `round` in the generated CSV. The 147/867 visit count is 16.955%, presented as 17.0%.

Review and quality limits

Both review-register lists cover all 1,003 admitted core IDs; 10 additional reviewed IDs are outside the admitted core. Coder IDs identify AI-assisted second coding and adjudication. This is not human peer review. No measured inter-coder agreement is available.

The row status remains pending final adjudication for 944 core records: 844 current, 100 historical. Completed-adjudication labels occur for 59: 23 current, 36 historical. Registry decision strings also retain the suffix `provisional`. Review-list coverage, row-status completion, and hash validation are different forms of evidence. The frozen summary's statement of complete second review cannot establish that every final adjudication was completed. The reanalysis preserves this discrepancy and labels the article preliminary.

Before treating the results as final: reconcile statuses with underlying decision/adjudication documents; review code evidence and heterogeneous scope labels; document agreement and correction counts. Reproducing frozen totals alone does not independently validate coding. Article examples are paraphrases from the coded registry. Stripe and Aiven sources were readable during the editorial review; Smarsh failed retrieval and Virgin Voyages returned no readable text. Those two cases are explicitly attributed to the frozen register.

Source availability and capture metadata are incomplete for some records. Posting dates are recorded for 29/867 current and 12/136 historical, with 962/1,003 missing overall. Collection rounds changed discovery coverage. No longitudinal, labor-market prevalence, hiring, or causal estimate is supported by this design. No confidence intervals for the wider market are reported. Salary, skills, seniority, geographic and business-model comparisons require additional normalization and coverage audits; raw heterogeneous labels are not used for those headline claims.

Data tables

Sample register, CSV

Outcome tables, CSV

References and source materials

1. **AI Search Jobs 2026: what 867 postings ask for, and what the evidence shows**
https://fernandoangulo.com/leadership/ai-search-jobs

2. **Methodology**
https://fernandoangulo.com/leadership/ai-search-jobs-methodology

3. **sample register**
https://fernandoangulo.com/downloads/ai-search-jobs-sample-2026-09-22.csv

4. **cohort tables**
https://fernandoangulo.com/downloads/ai-search-jobs-outcomes-2026-09-22.csv

5. **Stripe source**
https://stripe.com/careers/listing/head-of-aeo-seo/8128634

6. **Aiven posting**
https://aiven.io/careers/job/4838352101

7. **original source**
https://red.wd1.myworkdayjobs.com/en-US/VV/job/Sr-Discoverability--GEO-SEO--Manager_R2364

8. **original source**
https://jobs.lever.co/smarsh/9cf31fda-6d86-428c-b2c1-d09e03d2de8f

9. **topic ownership map**
https://fernandoangulo.com/leadership/topic-ownership

10. **reporting model**
https://fernandoangulo.com/leadership/geo-reporting