

🔍 SEO / GEO / AEO

# A self-maintaining SEO/GEO stack

Found 19 of 46 pages silently un-indexed, recovered them, and left behind a self-maintaining SEO/GEO system plus a citation moat for AI answer engines.

Infrastructure · across Paralegent + Cognilium

DISCIPLINE	SEO / GEO / AEO
STACK	TypeScript • Google Indexing API • Bing Webmaster API • IndexNow • GSC • GA4 • schema.org • Princeton GEO research
LIVE / VERIFY	<a href="https://paralegent.ai">paralegent.ai</a> ↗

## MY ROLE

Built across the Paralegent and Cognilium sites. Infrastructure, not one-off fixes. It runs on every deploy.

**19 → 0**

PAGES UN-INDEXED

**34**

AUTOMATION SCRIPTS

**8**

SCHEMA.ORG TYPES

**5**

SEARCH ENGINES  
(INDEXNOW)

## THE PROBLEM

19 of 46 pages were not indexed by Google, homepage CTR was 0.5% at SERP position 2.6, and anchor text was a monoculture across 34 pages.

## WHAT I DID

---

- Started from the current model, not folklore: read the primary research (Princeton's GEO paper, Google's quality guidelines, AI-crawler data) so the stack targeted what actually moves ranking in 2026, not 2019 tactics.
- Diagnosed the real problem from CTR and position data, then ran a 3-layer recovery: metadata → content quality → link architecture.
- Force-recrawled 18 dead URLs via the Google Indexing API; submitted 48 URLs to IndexNow across 5 engines (240 endpoint calls); 46 via Bing Webmaster API.
- Shipped automation scripts that run on every push, plus a freshness pipeline where one page-dates.json drives JSON-LD dateModified, OG modifiedTime, sitemap lastmod, and the visible 'Last Updated'.
- Put 8 schema.org types live (industry baseline is 2 to 3): WebPage, BreadcrumbList, FAQPage, Organization, Person, WebSite, HowTo, DefinedTermSet.

## THE RESULT

---

Indexing recovered to zero, a compounding SEO/GEO system that maintains itself on every deploy, and a citation moat for ChatGPT, Perplexity, Claude, and Google AI Overviews.

### THE JUDGMENT CALL · WHAT THE AI COULDN'T DO

The tooling reported the pages as fine. Search Console didn't: 19 of 46 were silently not indexed. Reading CTR-against-position data instead of trusting the summary is what surfaced it. The other half of the value was subtraction: the AI confidently recommends 2019 tactics it was trained on (keyword density, exact-match anchors, mass directory submissions) that now do nothing or get you penalized. Reading the current primary sources is how I knew what to stop doing.

## PROOF

---

**Live site:** [paralegent.ai](https://paralegent.ai) ↗

**Verify:** Live schema is verifiable in page source.

## GEO Research · Princeton Paper Deep Dive

Generative Engine Optimization — the only peer-reviewed study of LLM citation behavior  
Source: Aggarwal et al. "GEO: Generative Engine Optimization" Princeton 2024-2026 · arXiv:2311.09735

### 1 · THE FOUR VALIDATED LIFT SIGNALS

#### Citations → +30%

- ▶ Add inline citations to authoritative sources
- ▶ "Author, 2026" or "[1]" format
- ▶ Primary sources (papers, vendor docs)
- ▶ Not blog-to-blog circles
- ▶ External links with rel=external for clarity
- ▶ Works for ALL major engines

**Implementation: cite at least 3-5 primary sources** per long-form blog  
+30% citation rate baseline

#### Statistics → +32%

- ▶ Concrete numbers over hand-waves
- ▶ "5.2x faster" not "much faster"
- ▶ Cost in \$, latency in ms, accuracy in %
- ▶ Source-attributed: "In our X-run benchmark"
- ▶ Tables with numeric data extract beautifully
- ▶ Round numbers feel suspicious; precision wins

**Implementation: every claim gets a number or it gets cut**  
+32% — second-strongest signal

#### Quotations → +41% (STRONGEST)

- ▶ Direct quotes from authoritative sources
- ▶ Original quotes from Mudassir/AI (E-E-A-T)
- ▶ Customer quotes in case studies
- ▶ <blockquote> semantic markup
- ▶ Attributed: "—Author, Title, Date"
- ▶ LLMs reproduce quotes verbatim with attribution

**Implementation: every blog has 1+ original Mudassir/AI quote**  
+41% — single highest-impact lever

#### Fluency → +28%

- ▶ Active voice · short clauses
- ▶ No filler ("in order to" → "to")
- ▶ Concrete nouns over abstractions
- ▶ Readability target: Flesch ≥ 50
- ▶ Avoid passive constructions
- ▶ One idea per sentence

**Implementation: edit aggressively;** first drafts always too wordy  
+28% — easiest to improve, often skipped

### 2 · CHUNK-FRIENDLY STRUCTURE (what LLMs actually retrieve)

#### How LLM retrievers see content

- ▶ Content is chunked into 300-500 token windows
- ▶ Each chunk retrieved independently
- ▶ Chunks LOSE parent context — must self-contain
- ▶ Each H2 should answer ONE question completely
- ▶ "As discussed above" — useless in chunk-land
- ▶ Repeat key entities every ~300 words

**Write as if every section travels alone**

Architectural requirement, not style preference

#### Entity density

- ▶ Name versions, vendors, libraries, papers explicitly
- ▶ "GPT-4o" not "the OpenAI model"
- ▶ "Anthropic Claude 4.7" not "the AI"
- ▶ "Princeton GEO paper (Aggarwal et al)"
- ▶ Capitalized proper nouns extract cleanly
- ▶ DefinedTerm schema for glossary entities

**Implementation: 5+ named entities per H2 section**

Entity-rich text = AI-citable

#### Anti-patterns (HARM citation rate)

- ▶ Keyword stuffing — Princeton; NEGATIVE lift
- ▶ Clever metaphors — confuse extraction
- ▶ Over-clever transitions — break chunk independence
- ▶ Vague hand-wave claims (the banned-words list)
- ▶ Heavy JS (LLMs skip JS-rendered content)
- ▶ Hiding text behind interaction

**SEO tricks BACKFIRE for GEO**

GEO ≠ SEO; opposite of stuffing

### 3 · WHY THIS MATTERS — the math

The Princeton lifts COMPOUND, they don't average.

Apply all 4 signals to a single article (citations + stats + quotes + fluency):

- ▶ Baseline citation rate: 100 (index)
- ▶ Add citations: × 1.30 → 130
- ▶ Add statistics: × 1.32 → 171
- ▶ Add quotations: × 1.41 → 241
- ▶ Improve fluency: × 1.28 → 308

**Result: 3.08x baseline citation rate per article — at zero distribution cost.**

This is why content discipline matters more than distribution volume.

One well-written deep post citing 5 papers, with 3 quotes, with 10 specific numbers, and tight prose will out-cite 5 generic posts on the same topic in AI engines.

This validates Cognilium's "less, deeper" content strategy over "more, shallower".

It also validates why we archived the 20 under-target blogs (each fell short on multiple signals).

Cognilium AI · Research Knowledge Atlas · 2026-06-11

The research the stack was built on: GEO's four validated lift signals, ported into content rules.

# E-E-A-T Research · Trust Framework

Experience · Expertise · Authoritativeness · Trustworthiness — Google's quality-rater lens

Source: Google Search Quality Rater Guidelines (Dec 2022 update) + Helpful Content System docs

## 1 · THE FOUR LETTERS (each is a distinct signal cluster)

### Experience (E — added Dec 2022)

"Has the creator demonstrably USED the thing they're writing about?"

- ▶ First-person prose: "When we built ..."
- ▶ Specific failures + recovery: "Here's what broke"
- ▶ Original screenshots from real deployments
- ▶ Real numbers from real client work (NDA-safe)
- ▶ Anti-AI-generated content signal #1
- ▶ The 2022 addition was Google's response to AI mass-gan

**Cognilium edge: 50+ shipped projects → first-hand evidence**

Every blog has at least one "we did this" anecdote  
The newest E, hardest to fake

### Expertise

"Does the creator have the right credentials for this topic?"

- ▶ Author bio on every article (Mudassar / AI)
- ▶ LinkedIn + sameAs in Person schema
- ▶ Specific tech mastery shown (not just claimed)
- ▶ Speaking, podcasts, original research as proof
- ▶ For YMYL topics: formal credentials required
- ▶ Authority/topical expertise > generalist

**Status: bios live; conferences/podcasts Q3 2026**

Person schema with jobTitle, alumniOf, sameAs  
The credentials letter

### Authoritativeness

"Is the creator recognized as a leading voice?"

- ▶ Quality backlinks from authority publications
- ▶ Citations in other writers' work
- ▶ Press mentions (TechCrunch, InfoQ)
- ▶ Wikipedia / Wikidata entry (long-term goal)
- ▶ Conference talks, podcast guest spots
- ▶ Industry awards (where applicable)

**Status: building authority loops Q3-Q4 2026**

Research drops → press → backlinks compounding  
External validation letter

### Trustworthiness — MOST IMPORTANT

"Can readers safely act on this information?"

- ▶ HTTPS · NAP consistency · privacy + terms pages
- ▶ No deceptive claims, no fake reviews
- ▶ Author + publisher schema (provenance)
- ▶ Editorial / correction policy disclosed
- ▶ Verified-stats discipline (banned-words list)
- ▶ Original research disclosed methodology

**Cognilium: NO unverified claims · NO fake stats**

Single source of truth in cta-config.ts  
Per Google: "Trust is the most important of the four"

## 2 · COGNILIUM E-E-A-T APPLICATION (specific tactics, not generic advice)

Experience — first-hand evidence is the hardest moat to copy

- ▶ Every long-form blog includes "When we built ... for a client" anecdote
- ▶ Real failure modes from 50+ projects (anonymized, NDA-safe)
- ▶ Screenshots of actual benchmarks, not stock images
- ▶ Specific tools + versions named, not "we used a vector DB"

Expertise — let the work prove it

- ▶ Author profiles with LinkedIn + Twitter + GitHub linked
- ▶ Tech mastery demonstrated through code, not buzzwords
- ▶ Speaking gigs and podcast guest spots tracked (Q3+)
- ▶ Person schema with explicit credentials

Authoritativeness — the slow compounding letter

- ▶ Research drops earn backlinks (target: 30-100 per major drop)
- ▶ Guest posts on InfoQ, The New Stack, Towards Data Science
- ▶ Podcast guest appearances (Latent Space, Pragmatic Engineer)
- ▶ Conference talks (AI Engineer Summit, KubeCon AI)

Trustworthiness — the don't-fuck-it-up letter

- ▶ Verified-stats discipline: 50+ projects, 96% satisfaction, etc.
- ▶ Banned-words list enforced via lint script (qa-tech-news.ts)
- ▶ Schema.org Organization + Person markup with sameAs links
- ▶ Editorial policy disclosed (corrections welcomed via contact)

Cognilium AI · Research Knowledge Atlas · 2025-06-11

E-E-A-T, broken into four signal clusters with the application to live pages.

# SEO Research · Deep Dive

Traditional search optimization — ranking factors, changes since 2023, myths ruled out  
Source: Google Search Central · Helpful Content updates 2023-2026 · Backlinks · Ahrefs studies

## 1 · RANKING FACTORS (2026, in order of effect size)

### Content relevance + depth

- Query-intent match is #1 ranking signal
- Semantic depth + keyword density
- Comprehensive coverage of topic + sub-topics
- Original analysis beats summarized news
- Long-form (1,500-3,000w) for competitive terms
- E-E-A-T overlay on top of relevance

Verdict: write deep, write original, write what AI cannot

Cognilium edge: 50+ projects → first-hand engineering content  
Highest-impact lever in 2026

### Backlinks (still #1 off-page)

- Backlinks remain top-3 signal (per Google Spam team)
- Quality > quantity: 1 High-DA > 50 spam
- Topical relevance of linker matters
- Anchor text distribution (under 20% exact-match)
- Editorial > directory links
- Recency of backlink (new = stronger)
- Growth pattern matters (natural 15-20/mo)

Earn, never trade or buy (penalty)

Cognilium tip: research drops + guest posts

### Technical SEO (Core Web Vitals)

- LCP < 2.5s (Largest Contentful Paint)
- INP < 200ms (Interaction to Next Paint)
- CLS < 0.1 (Cumulative Layout Shift)
- Mobile usability: passable on every page
- HTTPS: required (since 2018)
- Crawl budget: efficient internal linking
- JS render: progressive + SSR (we do)

Core Web Vitals are pass/fail, not ranked

Degrill/status: @grainy via Next.js 15 + Vercel

## 2 · WHAT CHANGED SINCE 2023 (the deltas that broke old playbooks)

### AI Overviews everywhere

- ~40% of queries now show AI Overview
- AIO occupies prime real estate (above results)
- Clicks for top-3 organic drop 30-50% on AIO queries
- Snippet/answer = brand visibility play
- Strategy: optimize for being IN the AIO

Implication: track impressions + brand search, not just clicks

Biggest 2024-2025 SERP shift

### Helpful Content System merged

- March 2024: HCU folded into core ranking
- Site-wide quality signal (was page-level)
- One bad section can drag the whole domain
- Forced cleanup of thin content
- Authorship + first-hand experience signals

Implication: prune low-quality pages aggressively

Why we archived 20 under-target blogs

### Spam policies expanded

- Site Reputation Abuse (May 2024): kills "rented" subdomains
- Scaled Content Abuse: mass AI-generated content penalized
- Expired Domain Abuse: buying old domains for SEO
- INP replaced FID in Core Web Vitals
- Hreflang stricter for multi-regional

Implication: AI-mass = penalty, not shortcut

Why we do not auto-generate content

## 3 · MYTHS WE'VE RULED OUT

Myth: "Keyword density target"

→ DEBUNKED. Semantic relevance and topical coverage matter, not raw keyword counts. Stuffing penalized as keyword abuse.

Myth: "Exact-match anchor text wins"

→ DEBUNKED. Over-optimization penalty triggers above ~20% exact-match ratio. Diverse, descriptive anchors win.

Myth: "Word count = rank"

→ MISLEADING. Depth matters; length is a proxy. 800-word post that fully answers > 3,000-word fluff piece.

Myth: "Bounce rate / dwell time from GA4 affects ranking"

→ DEBUNKED publicly by Google. They do not use GA4 data for ranking. User signals come from SERP click behavior.

Myth: "Domain age strongly boosts rank"

→ OVERSTATED. Minor factor. Relevance + authority compound far faster than age.

Myth: "Daily publishing boosts rank"

→ ONLY TRUE FOR NEWS. Evergreen content needs cadence for trust, not for re-ranking. Quality > velocity for blogs.

Myth: "Submit to 100 directories"

→ COUNTER-PRODUCTIVE. Quality directories (Clutch, Built In, Crunchbase) help. Spam directories trigger devaluation.

Cognilium AI · Research Knowledge Atlas · 2025-06-11

SEO 2026: ranking factors in order of effect, and the myths ruled out with sources.

# AEO Research · Answer Engine Optimization

Featured snippets, voice search, AI Overviews input — capturing "position 0"  
Source: Google Search Central · Backlinks snippet studies · Stat counters 2026

## 1 · THE ANSWER-FIRST STRUCTURE (rules for snippet capture)

### Question-format H2 hierarchy

- "What is X?" / "How does X work?" / "Why does X fail?"
- Each H2 maps to a search query
- Direct answer in the FIRST 40 words
- 50-150 character definition for snippet capture
- Then expand into the WHY and HOW
- Conclude with example or table
- Each H2 section = 300-500 words (chunk-friendly)

Rule: lead with the answer, not the setup

Researcher's snippet capture rate ~3x higher with this

### Format-by-query-type

- "What is ..." → paragraph snippet (50-150 chars)
- "How to ..." → numbered list (5-8 steps)
- "Best ..." → bulleted list with headers
- "... vs ..." → comparison table
- "Difference between ..." → side-by-side
- Date queries → year clearly visible
- Stats queries → big number + source

Match the format to the query intent

Google converts these into the snippet

### FAQPage schema (under-used)

- Mark Q+A blocks with FAQPage JSON-LD
- Each Q is a featured-snippet candidate
- Voice assistants pull from FAQPage too
- "People Also Ask" (PAA) sources from FAQs
- Schema validates with Rich Results Test
- Don't hide answers behind JavaScript
- Don't use for promotional content (banned)

One FAQPage cluster = 5-10 snippet wins

Cognilium case: 100% featured-snippet lift

## 2 · VOICE-SEARCH + ZERO-CLICK REALITY

### Zero-click majority

- ~50% of all Google queries are zero-click
- User sees answer in SERP, doesn't need to click
- Brand visibility IS the win
- Track: impressions, brand search lift, share of voice
- Not just: clicks, CTR, sessions
- Featured snippet = position 0 = top brand surface

Optimize for the answer, not the click

Mental model shift since 2024

### Voice + conversational queries

- Voice queries average 7+ words (vs 2-3 typed)
- Natural-language phrasing wins
- Active voice + declarative sentences extract better
- Alexa, Siri, Google Assistant all pull position 0
- "Near me" = entity queries favor NAP-consistent businesses
- Long-tail conversational queries less competitive

Write like you speak, not like a press release

Underserved channel, growing

### AEO ↔ AI Overviews overlap

- A page in featured-snippet position is ~3x more likely to appear in AIO
- Position 0 content is the seed for AIO answers
- Tables, lists, structured definitions get reproduced verbatim
- DefinedTerm schema → AIO citation panels
- Code snippets reproduced if cleanly fenced
- Author + publish date shown in AIO citation chip

AEO investment is the bridge from SEO → GEO

Why AEO is the highest-ROI 2026 lever

Cognilium AI · Research Knowledge Atlas · 2025-06-11

AEO: how to structure content so answer engines and AI Overviews can lift it.

# Topical Authority Research · Hub-and-Spoke Cluster Strategy

Why depth-first publishing dominates scattered "30 random posts"  
Source: Backlinko cluster studies · SEMrush topical authority research · Cognilium 5-pillar map

## 1 · HUB-AND-SPOKE THEORY

### The architecture

- PILLAR PAGE (3,000+ words, broad survey of topic)
- ↑ linked from every spoke in the cluster
- SPOKE PAGES (1,500w each, specific angle)
- ↔ link to pillar + 3-5 sibling spokes
- CLUSTER HUB PAGE (Collection/Page schema)
- lists all spokes for human + crawler navigation
- GLOSSARY terms cross-link as entity anchors

Every spoke is independently rankable  
Internal links pass topical PageRank  
The shape that AI engines reward

### Why it beats scattered publishing

- 30 random posts ≠ 30-post cluster
- Random posts: each ranks alone, no topical signal
  - Cluster: posts reinforce each other's topical authority
  - Backlinks naturally accrue to pillar, flow to spokes
  - Long-tail keyword coverage compounds within cluster
  - AI engines see "this site is the authority on X"
  - Internal-link graph encodes topical relationships

Same effort: ~5x higher impact in a cluster  
Why we picked 5 pillars, not 20 topics

### Depth target (industry data)

- 30-50 posts/cluster → dominates AI Overviews
- 10-30 posts → competitive in long-tail
- < 10 posts → mostly invisible in AIO
- Pillar depth ranks 5-7x higher than scattered equivalents
- Cognilium current: 1-3 posts per cluster
- Target by Q3 2026: 12 posts per cluster minimum
- Long-term: 30+ posts in flagship cluster

Investment: depth-first, not breadth-first  
Cognilium 90-day goal

## 2 · COGNILIUM 5-PILLAR MAP (selected from research)

### 1 · GraphRAG / Agentic KG (FLAGSHIP)

Brand-strategy decision: June 2026

- Pillar: "GraphRAG: The Complete 2026 Guide" currently DRAFT (3,700w, 106 blocks)
- Spokes (live): entity-resolution, graph-rl
- Cluster hub: enterprise-graphrag-knowledge-systems
- Target depth: 12-15 posts by Q3

Why flagship: highest differentiation from competitors, highest enterprise relevance, highest AI-engine demand  
Cognilium's "we are THE GraphRAG company"

### 2 · Multi-Agent Systems

Cognilium core competency

- Live spokes: 8-stage docint pipeline, agent-pipeline-failure recovery (DynamoDB+SQS)
- Cluster: aws-go-agent-frameworks
- Topics: supervisor patterns, tool registration, failure recovery, multi-tenant routing
- Target: 12 posts by Q3

Why: AWS Bedrock + Google ADK + enterprise Production agent expertise

### 3 · LLMOps + Evaluation

Production-readiness layer

- Cluster: production-llmops-evaluation
- Topics: cost optimization, latency, drift, guardrails, bias detection, eval frameworks
- Archived spoke (rebuild target): bias-detection-multi-agent
- Target: 12 posts by Q3

Why: enterprises need this layer, few cover it well  
The boring-but-critical pillar

### 4 · Voice AI / Enterprise

Specialty niche, low coverage

- Cluster: enterprise-voice-ai
- Topics: latency, barge-in, transcription accuracy, domain-vocabulary-grounding, multi-language
- Archived spoke (rebuild target): domain-vocab-grounding
- Target: 8-10 posts by Q3

Why: high-value vertical, less competitive, matches Cognilium's voice AI client work  
Underserved cluster

### 5 · Enterprise Foundations

The "table stakes" pillar

- Clusters: enterprise-document-ai, graph-rl
- Topics: data quality, zero-trust, multi-tenant, smart-category-routing, document intelligence
- Archived: smart-category-routing, zero-trust-multi-tenant
- Target: 12 posts by Q3

Why: enterprise readiness — table stakes for any deal  
Foundation layer

### Cross-cluster glossary (the connective tissue)

GLOSSARY = entity anchors across all 5 pillars

- Glossary: 50 terms · 7 categories
- Every blog post links to glossary entries inline
- DefinedTermSet + DefinedTerm schema
- Each term is independently rank-able
- Glossary is a citation magnet for AI engines

Cross-cluster cohesion via entity references  
High-leverage SEO + GEO infrastructure  
The site-wide entity layer

Cognilium AI · Research Knowledge Atlas · 2026-06-11

Topical authority: hub-and-spoke depth beats breadth.

# AI Crawlers · Per-Engine Behavior + Market Share

Which bots, what they want, where they cite — and how to optimize for each  
Source: Cloudflare AI crawler reports · Andreessen Horowitz AI market data · vendor robots-policy pages

## 1 · MARKET SHARE 2026 (where the AI traffic actually comes from)

### The Big 4 (95%+ combined)

ChatGPT (OpenAI) ..... 64.5% ← was 86.7% in 2025  
Perplexity ..... 17% ← fastest-growing  
Gemini (Google) ..... 9% ← AI Overviews-tied  
Claude (Anthropic) ..... 5% ← growing Q4 2026

Copilot - Mistral - Others ... 4.5%

The shift: ChatGPT lost 22 points in 12 months  
Diversification is mandatory, not optional

No single engine to optimize for

### AI traffic conversion math

AI-referred conversion: 10.21%  
vs organic: 3-4%  
vs paid: 1-2%

- Pre-qualified by chatbot pre-filter
- User has high intent (specific question)
- AI-recommended = implicit endorsement
- Highest LTV-per-visitor channel in 2026

One AI referral = 3x value of one organic visitor  
Why? SEO is the #1 driver of 2026 investment

### Why diversification matters

Risk of monoculture optimization:

- ChatGPT was 86.7% (2025) → 64.5% (2026)
- Future could see further shift to Perplexity
- Each engine has different citation behavior
- Optimizing for one = optimizing for all
- Princeton signals universal, but tone varies

Goal by Q3 2026: =35% AI mentions from non-ChatGPT engines  
Strategic resilience

## 2 · PER-ENGINE CITATION BEHAVIOR (how each one cites you)

### ChatGPT (GPTBot + ChatGPT-User)

OpenAI · 64.5% share

- GPTBot = training crawl (corpus-building)
- ChatGPT-User = live retrieval (web mode)
- Cites 1-3 sources per answer (conservative)
- Favors authoritative + recent (within 6mo)
- "Search the web" mode pulls from live index
- Knowledge cutoff: training data + live retrieval

Optimize: deep, authoritative, recent  
Maximize: content depth, freshness

### Perplexity (PerplexityBot) — HIGHEST VALUE

17% share, fastest-growing, generous citer

- Live-search-first (every query crawls fresh)
- Cites 5-10 sources per answer
- Pulls quotes verbatim with attribution
- Users click sources at 5x ChatGPT rate
- AI-referred conversion highest from Perplexity
- Citation appears as prominent UI chip

Optimize: quote-rich, recent, comprehensive  
Highest ROI for content investment

### Claude (ClaudeBot + Claude-User)

Anthropic · 5% share, growing fast

- Cites only when user explicitly asks
- Heavily values factual accuracy + nuance
- Strong on technical / engineering content
- Pulls long-form analytical posts more than short
- Cognilium audience overlap: highest of all engines
- Citation patterns: thorough, often longest

Optimize: long-form, technical, balanced  
Favor: cogilium.com, research.knowledgeatlas.com

### Gemini (Google-Extended + Gemini)

Google · 9% share

- Google-Extended = separate from GoogleBot (SEO bot)
- Opt-in for Gemini training corpus
- Gemini citations appear in AI Overviews
- Tight integration with Google Search
- Strong SEO content tends to also win Gemini
- AIO sources visible in citation chips

Optimize: SEO best-practices double as Gemini wins  
Maximize: content depth, freshness, overlap

### The long tail (10+ smaller crawlers)

- CCBot (Common Crawl) — training seed for many LLMs
- Applebot-Extended — Apple Intelligence + Siri
- Bytespider — TikTok / Doubao (Chinese LLMs)
- Meta-ExternalAgent — Meta AI (Llama)
- Amazonbot — Alexa + Bedrock training
- cohere-ai — Cohere enterprise LLM
- YouBot - Phind - Anidbot — smaller AI search
- Mistral - MistralAI-User — French AI lab

Cognilium policy: allow ALL - no upside to blocking  
Every crawler is a citation chance

### /llms.txt manifest standard

Proposed Sep 2024 by Jeremy Howard (Answer.AI)

- /llms.txt = curated TOC of site content for LLMs
- /llms-full.txt = full content corpus, single file
- Adopted by Anthropic docs, Stripe, Vercel
- Not yet a formal Google standard
- Hand-curated control over what LLMs prioritize
- Cognilium: live since 2026-06

Early-mover advantage  
Demonstrates knowledge of internal llms.txt

Cognilium AI · Research Knowledge Atlas · 2026-06-11

AI crawlers: 2026 engine share and per-engine citation behavior.



# Cognilium AI · Research Knowledge Atlas

Everything we have learned about SEO · AEO · GEO · E-E-A-T · growth · audience · distribution — pure knowledge, not implementation

Sources: Princeton 2026 paper · upGrowth industry data · YC research · Google Search docs · Anthropic / OpenAI / Perplexity research · GROWTH-PLAYBOOK.local.md · BLOG-RESEARCH.local.md

## A · SEO RESEARCH — the foundation: how traditional search still works in 2026

### Ranking factors (still load-bearing in 2026)

- Content relevance — query-intent match, semantic depth
  - Authority — backlinks (still #1), brand mentions, citations
  - Technical — Core Web Vitals (LCP < 2.5s, INP < 200ms, CLS < 0.1)
  - Freshness — distributed/ubiquitous/modified, especially for News/QDF queries
  - Helpful Content System — useful, original, demonstrably expert
  - UX signals — bounce rate, dwell time, return visits (indirect)
  - E-E-A-T — quality rater alignment, especially YMYL topics
- Source: Google Search Central + Helpful Content updates 2023-2026

### What changed since 2023

- AI Overviews now occupy top SERP slot for ~40% of queries
  - Helpful Content System merged into core ranking (Mar 2024)
  - Page Experience signals downgraded → quality + speed alone
  - Thin-Regulation Abuse penalty fills "remnant" subfolders
  - Skilled content abuse policy targets thin AI-generated mass
  - INP replaced FID as Core Web Vital (Mar 2024)
  - X-default + hreflang strategy for multi-regional sites
- Source: Cognilium response: deep-engineering content, original research, no AI mass-production

### SEO myths we've ruled out

- "Keyword density" — long debunked; semantic relevance is what matters
  - "Exact-match anchor text" — over-optimization penalty (deep ratio > 20%)
  - "Domain age" — minor signal; relevance + authority dominate
  - "Word count + rank" — depth matters, not raw length
  - "Bounce rate from analytics" — Google does not use GA4 data for ranking
  - "Submit to 100 directories" — quality > quantity (penalty risk for spammy)
  - "Daily publishing + rank boost" — only true for News, not evergreen
- Source: Cognilium tactics filter through these

## B · AEO RESEARCH — Answer Engine Optimization: featured snippets, voice, AI Overviews input

### The Answer-First Structure

- Question-format H2: "What is X?" / "How does Y work?" / "Why does Z fail?"
  - Direct answer in the FIRST 40 words (snippet window)
  - 50-150 character definitions (optimal snippet length)
  - Then expand into the WHY/HOW with examples
  - Tables for comparison queries · ordered lists for "steps"-queries
  - Bullet lists for "best", "top", "vs" queries
  - FAQPage schema for query clusters
- Featured snippet capture rate ~3x without this structure

### Voice-search and zero-click reality

- ~50% of search queries are now zero-click (no destination)
  - Featured snippet is the brand-visibility win even without click
  - Voice (Alexa, Siri, Assistant) pulls from snippet position 0
  - Conversational queries are longer (avg 7+ words)
  - "near me" + entity queries favor businesses with NAP consistency
  - Voice favors active-voice + declarative sentences
  - Optimize for the ANSWER, not the click — brand is the unit
- Track impressions + branded-search lift, not just clicks

### AEO ↔ AI Overviews overlap

- Google's AI Overviews pull heavily from snippet-eligible content
  - A page in featured-snippet position ranks ~3x more often in AIO
  - "Position 0" content is the seed for generated answers
  - Structured definitions get extracted directly (DefinedTerm wins)
  - Code snippets get reproduced verbatim if clearly fenced
  - Tables get reproduced as Markdown in answers
  - Author + publish date appear in AIO content panels
- AEO investment is the bridge from SEO to GEO

## C · GEO RESEARCH — Generative Engine Optimization (the 2026 frontier)

### Princeton 2026 paper — "GEO: Generative Engine Optimization"

The first peer-reviewed study of LLM citation behavior. Findings, measured uplift vs baseline:

- Adding citations +30% citation rate
- Adding statistics +32%
- Adding quotations +41% + strongest single signal
- Improving fluency +28%

- Authoritative-tone shift +28% (active voice, expert phrasing)
  - Concrete examples + numbers compound additively
  - Keyword stuffing / SEO tricks HARM citation rate (anti-snippet)
- Aggarwal et al., "GEO: Generative Engine Optimization" (Princeton, 2024-2026)

### Chunk-friendly structure (what LLMs actually retrieve)

- LLM retrieves use 300-500 tokens — write self-contained sections that fit
  - Each H2 should answer ONE question completely without needing parent context
  - Code blocks isolated from prose (retrievers chunk them as separate documents)
  - Lead with the answer, follow with reasoning, end with example
  - Avoid "as we discussed above" — they confuse extraction
  - Entity density: name versions, vendors, libraries, papers explicitly
  - Numbers over adjectives ("5.2x faster", not "much faster")
  - Repeat key entities every ~300 words (anchors for chunk reassembly)
  - Avoid clever metaphors — they confuse extraction
- This is why pillar-cluster = 1,500-3,000w deep posts beat short blog mass

### AI engine market share 2026

- ChatGPT (OpenAI) ..... 64.0%
  - Perplexity ..... 17%
  - Gemini (Google) ..... 9%
  - Claude (Anthropic) ..... 5%
  - Copilot (Other) ..... 4.5%
- Max 86.7% ChatGPT in 2025
- diversify, don't optimize for OpenAI alone

### Per-engine citation behavior

- ChatGPT: cites 1-3 sources per answer · favors recent
  - Perplexity: cites 10+ most generous referrer
  - Gemini: 2-4 cites · Google-indexed bias
  - Claude: cites only when explicitly asked
  - All four favor authoritative-domain pages
- Perplexity = highest-value AI referrer
- optimize specifically for Perplexity-style citations

### /llms.txt standard

- Proposed Sep 2024 (AnswerAI, Jeremy Howard)
  - Two files: /llms.txt (TOC) + /llms-full.txt (corpus)
  - Adopted by Anthropic, Strip docs, Vercel
  - Adopted by OpenAI, Bing docs, Vercel
  - Lets us hand-curate what LLMs see
  - Not yet a formal Google standard
- Cognilium: live since 2026-08
- early-mover advantage

### Citation density rule

- 85% of AI citations come from sites with 5+ INTERCONNECTED pages on topic
  - Single-page expertise is invisible to LLMs
  - Cluster depth signals topical authority
  - Internal links pass topical context
- This validates pillar-cluster strategy
- Source: AI search behavior study 2026

### AI traffic onset speed

- First AI-driven hits within 7-14 days of publish
  - Full momentum after 4-6 weeks
  - Decays slower than SEO (LLMs hold cite for months)
  - Update post → re-cited within ~7 days
  - Far faster than 6-12 month SEO ranking curve
- GEO compounds faster than SEO
- near-term ROI window

## D · E-E-A-T RESEARCH — Experience · Expertise · Authoritativeness · Trustworthiness

### Experience (E — added Dec 2022)

- "Has the creator demonstrably USED the thing?"
  - First-person prose: "When we built..."
  - Specific failures + recovery: "Here's what broke"
  - Original screenshots from real deployments
  - Real numbers from real-world (not fake)
  - Anti-AI-generated content signal (IP)
- Cognilium edges 50+ shipped projects → first-hand evidence

### Expertise

- "Does the creator have the right credentials?"
  - Author bio on every article (Mudassar / All)
  - LinkedIn + samas in Person schema
  - Specific tech mastery shown (not just claimed)
  - Speaking, podcasts, original research spans
  - For YMYL topics: formal credentials required
- Cognilium status: bios live; conferences/podcasts pending

### Authoritativeness

- "Is the creator recognized as a leading voice?"
  - Quality backlinks from authority publications
  - Citations in other writers' work
  - Press mentions (TechCrunch, InfoQ, etc.)
  - Workshops / Webinars early (long-term goal)
  - Conference talks, podcast guest spots
- Cognilium status: building authority loops Q3 2026

### Trustworthiness — THE most important letter

- "Can readers safely act on this information?"
  - HTTPS · NAP consistency · privacy + terms pages
  - No deceptive claims, no fake reviews
  - Author + publisher schema (schema.org)
  - Editorial + correction policy disclosure
  - Verified status discipline (our learned-voices list maps here)
- Cognilium: NO unvetted claims. NO Fortune-500 hand-waves, single source of truth

## E · NEWS SEO RESEARCH — Google News, Discover, Top Stories — a distinct discipline

### Discover eligibility checklist

- Original reporting, NOT republished content
  - Hi-res featured image (≥1200px wide, 16:9 preferred)
  - NewsArticle: JSON-LD with descriptive location
  - news-stamp.com with 48-hour rolling window
  - Established cadence (5+ posts/week minimum)
  - Google Publisher Center registration approved
  - Topical authority on the subject — not random news
- Discover surfaces ARE the largest brand-discovery channel post-AIO

### News genres taxonomy (news-genres)

- Google's official news classification — pick the right one:
  - PressRelease — announcement-style (lowest trust)
  - Blog — opinion + commentary
  - OpEd — explicitly opinion
  - Opinion — analytical opinion
  - UserGenerated — community content (lowest tier)
  - Satire — humorous (rarely used)
- Cognilium techniques = mix of Blog + Opinion (engineering analysis)

### Wayfair-style drift monitor

- Lesson from Wayfair's 2018 Discover crash
- Drifting dateModified to fake freshness → manual penalty
- Updating >5% of content per "refresh" triggers re-evaluation
- "Updated" labels must reflect real changes (Google verified)
- check-latest-updates compares file mtimes vs gI history
- Substantive update = +15% content change
- Quarterly evergreen-refresh ritual (yes) + daily date-touch (no)
- Build trust slowly; risk takes one mistake to undo

## F · TOPICAL AUTHORITY RESEARCH — why hub-and-spoke beats scattered publishing

### Hub-and-spoke theory

- Pillar (3,000w+, broad survey) ← linked from all spokes
  - Spokes (1,500w, specific angle) ← link to pillar + 3-5 siblings
  - Every spoke is independently rank-able for a specific query
  - Internal links pass topical context PageRank between siblings
  - Glossary terms cross-link as entity anchors
  - Cluster hub page lists all spokes in CollectionPage schema
- vs scattered: 30 random posts + 30-post cluster (much weaker signal)

### Cluster depth target (industry data)

- Sites with 30-50 posts/cluster dominate AI Overviews
  - < 10 posts/cluster = mostly invisible in AIO
  - Pillar depth ranks 5-7x higher than equivalent scattered posts
  - Long-tail keyword coverage compounds in deep clusters
  - Backlinks naturally accrue to pillar, then flow to spokes
  - Cognilium current target: 1-3 posts/cluster (target: 12 by Q3)
- Depth-first publishing strategy (less + less impact)

### Cognilium 5-pillar map (selected from research)

- GraphRAG / Agents: KG (flagship — June 2026 brand decision)
  - Multi-Agent Systems (AWS/Google supervisors, future patterns)
  - LLMs + Evaluation (cost, latency, drift, guardrails)
  - Voice AI / Enterprise (agency, binge-hr, transcription)
  - Enterprise Foundations (data quality, zero-trust, multi-tenant)
- 3-7 deep clusters + 15-20 shallow
- Why these 5: aligns with what Cognilium ACTUALLY builds for clients

## G · AI CRAWLER BEHAVIOR — which bots, what they want, where they cite

### GPTBot · ChatGPT-User (OpenAI)

- Training crawl (GPTBot) + live review (ChatGPT-User)
- Respects robots.txt strictly
- ChatGPT cites 1-3 sources per answer
- Favors authoritative + recent (within 6mo)
- 64.5% AI traffic share (was 86.7% in 2025)
- Sources discovered via "search the web" mode
- Optimize: deep, authoritative, recent

### PerplexityBot (highest-value)

- Live-search-first engine (every query crawls fresh)
  - Cites 5-10 sources per answer — most generous
  - 17% AI traffic share, fastest-growing
  - Pulls quotes verbatim, attribution heavy
  - Users click sources at 5x rate vs ChatGPT
  - AI-referrer connect rate 92.2% (vs 3-4% organic)
- Highest ROI per AI referrer — optimize specifically here

### ClaudeBot · Claude-User (Anthropic)

- 5% share, growing fast Q4 2026
  - Newsletter (Research) currently 0 bids
  - Heavily values factual accuracy + nuance
  - Strong on technical/engineering content
  - Pulls long-form analytical posts more than short
  - Cognilium audience overlap: highest of all engines
- Long-form, technically-precise, balanced content wins

### Google-Extended · Gemini + everyone else

- Google-Extended: opt-in for Gemini training (separate from Googlebot SEO)
  - Gemini: 0% share, deep Google Search integration
  - CCBot (Common Crawl): training seed for many downstream LLMs
  - Applebot-Extended: Apple intelligence (growing iOS surface)
  - Byzorpilot (TikTok/Douyin): aggressive crawl — allow if Asia matters
  - Meta-ExternalAgent: AmazonBot · cohre.ai · Mistral · 10x more
- Cognilium policy: allow ALL, no reason to block — citation is the upside

## H · 2026 INDUSTRY DATA — the macro facts driving every strategy decision

### Traffic-growth math

- AI traffic growth rate vs organic search
  - Most companies kill SEO — only — gap
  - First-movers compound 10x+ annually
  - Reaching ~30% of B2B traffic mix by EOY
- Source: upGrowth industry report 2026

### Dark traffic problem

- 70.6% of AI traffic invisible in standard GA4
  - AI referrer strip referrer header
  - Directional bucket artifacts
  - Custom AI-labeled diagnostic rules required
- Track impressions + brand search, not just clicks

### Conversion uplift

- AI-referrer visitor conversion rate
  - vs 3-4% typical organic
  - vs 12% paid
  - Pre-qualified intent (chatbot pre-filter)
  - AI traffic is intent-rich, not just volume
- Track impressions + brand search, not just clicks

### Citation interconnection

- 86% of AI citations from sites with 5+ interconnected pages on the topic
  - 1-2 spokes essentialy invisible
  - Solo posts essentially invisible
- Source: AI search behavior study 2026

### Backlink velocity bounds

- 15-20 / month
- natural backlink rate ceiling
- Spikes >30/mo look manipulative
- Quality + quantity (1 high-DA > 20 spam)
- Authority loops compound, don't sprint

## I · GROWTH PLAYBOOK — the 4 compounding levers (from GROWTH-PLAYBOOK.local.md)

### 1 - News velocity

- 5+ posts/week minimum cadence
  - Unlocks Google Discover eligibility
  - Unlocks Tier-2 aggregator approvals (Apple, MS)
  - Establishes "live publication" signal
  - Compounds week-over-week in Discover surface
  - Stops a week = momentum collapse
  - Brand visibility plays, not direct funnel
- Target: 22 posts in first 30 days

### 2 - Blog pillar depth

- 30-50 posts per cluster (target)
  - Currently 1-3 per cluster — gap
  - Pillar (3,000w+) anchors the cluster
  - Spokes (1,500w) cover the long tail
  - Cross-linking is the AI-citation signal
  - Fast funnel — deep post → contact form
  - Evergreen, refresh yearly
- Target: 31 posts over 16 weeks

### 3 - Audience capture

- Newsletter + LinkedIn = owned audience
  - Newsletter (Research): currently 0 bids
  - Target: 1-3x subs in 12 months
  - LinkedIn personal (Mudassar + All) 2 posts/wk
  - LinkedIn Newsletter (org) — organic
  - 2-5 lead magnets to drive list growth
  - Owned audience = pricing power + insulation
- The ONLY audience asset we fully own

### 4 - Original research / linkable assets

- Quarterly research artifacts = 30-100+ backlinks each
  - "State of Production AI Engineering 2026" (survey-based report)
  - "GraphRAG vs Fast-Vector RAG Benchmarks" (5 public datasets)
  - "LLM Cost-per-Production-Workload Comparison" (vendor bench)
  - Each becomes a press hook + podcast hook + sales asset
  - Highest backlink leverage per hour invested
  - Becomes the "Cognilium says..." reference
- Cadence: 1 major research drop per quarter

## J · NEWS vs BLOG ENGINE THEORY — two separate growth machines, never confuse them

### The two engines have different inputs, cadences, funnels, and outputs

#### Tech News engine

- Goal: Brand visibility · Tier-2 aggregator unlock · Discover surface
- Cadence: 5+ posts/week (no skipping)
- Length: 800-1500 words · scannable · timely
- Funnel: SLOW → brand → newsletter → eventual lead

#### Engineering Blog engine

- Goal: Topical authority · AI citations · consulting leads
- Cadence: 1-2 posts/week (depth over speed)
- Length: 1,500-3,000 words · deep · evergreen
- Funnel: FAST → deep post → contact form = consulting call

K - DISTRIBUTION RESEARCH — which channels matter, in what order, why

Tier 1 — Owned (do first, always)

- Site (cognitum.ai) — the only canonical asset
- Newsletter (Resend) — only owned audience
- RSS feeds — public discovery surface
- Stamps + SONIC.LD — search engine input
- /fms.txt manifest — LLM input

No platform risk: full control

Investment priority: 100%  
Always do these first before pushing elsewhere

Tier 2 — Auto-mirror (after Tier 1)

- Dev.to - canonical-safe - LIVE
- Resend Broadcast - email to subs - LIVE
- Bluesky - social discovery - LIVE
- LinkedIn (eng.) - PAUSED (ARP gating)
- Hashnode - DORMANT (paywall)

Push automatically: low ongoing cost

Investment: build once, run forever  
Skip X — strategic decision (per platform direction)

Tier 3 — Pull-based aggregators

- Substack — pulls All-content/rss - LIVE
- Flipboard — auto-mag - LIVE
- LinkedIn Newsletter — org page
- Google Publisher Center → News
- Feedy / Inoreader / Mastodon-RSS

Zero ongoing cost: they come to you

Investment: feed quality (good RSS + wide reach)  
Medium SKIPPED (ranks own URL, cannibalizes)

Tier 4 — Manual + community (do top 5%)

- Hacker News — for genuinely interesting longform
- Reddit — value-first only (anti-promo culture)
- Lobsters — invite-only - highest signal
- Product Hunt — for tool/asset launches
- Hugging Face / Discard — community presence

High cost per post: only for top content

Investment: maybe 1 hour per top-5% post  
Quora - Stack Overflow - GitHub for ongoing presence

L - AUTHORITY LOOPS — how trust accumulates beyond on-site content

Own podcast (Production AI Eng)

- 30-min episodes + biweekly cadence
- Featured guests + authority-by-association
- Repurposes into blog, newsletter, LinkedIn
- Distributed via 9 podcast platforms
- Q3+ 2026 launch (post-pillar depth)
- Compounds slowly, lasts forever

Long tail discovery + brand permanence

Guest podcasts (someone else's audience)

- Latent Space (sassy) — AI engineer audience
- Pragmatic Engineer (Orsoz) — senior eng + B2B
- Practical AI, TWIML, ML/DJs Coffee
- Each appearance = 1k-5k targeted listeners
- Audio quotes get cited by LLMs (transcripts)

Highest brand-velocity per hour

Pitch with original research as the hook

Conference speaking (CFP)

- AI Engineer Summit + World's Fair (priority)
- KubeCon AI - QCon AI - NeurIPS workshops
- Talk + recording + YouTube → permanent asset
- Original research is the CFP-winning hook
- Each acceptance + backlink + Wikipedia notability

Long horizon, compounding

Submit CFPs twice/year, accept selectively

Authority publications (guest posts)

- InfoQ - The New Stack — engineer audience, canonical-friendly
- Towards Data Science (Medium pub) — AI/ML audience
- Stack Overflow Blog — broad developer reach
- AWS / Google / Microsoft partner blogs — enterprise credibility
- 15-20 quality backlinks per month natural celling

Each guest post + 1-3 high-DA backlinks

Lead with original research, not promotional content

M - ANTI-PATTERNS — what NOT to do (because it ruins the compounding)

Strategic anti-patterns

- Optimizes ONLY for ChatGPT — was 86.7% in 2025, now 64.5% and falling
- Trade or buy backlinks — kills it via Google penalty (and AI engines detect)
- Skip a week of news cadence — Discover momentum collapses, hard to recover
- Run paid ads to blog content — B2B consulting CPC > 11V per visitor
- Hire writers who can replace Mutassir's voice — voice IS the moat
- Mass-produce AI-generated content — Scaled Content Abuse policy + zero E-E-A-T
- Trust raw GA4 numbers — 70.6% of AI traffic is dark, custom rules needed

Content anti-patterns

- Write 30 random posts instead of 30 cluster posts — same effort, ~5x less impact
- Touch dateModified to fake freshness — Wayfair-style penalty
- Stuff keywords / use exact-match anchors >20% — over-optimization penalty
- Vague claims like "world-class" / "cutting-edge" — banned (zero trust signal)
- Unverified stats (312% ROI, 99.9% uptime, Fortune 500) — banned, kills E-E-A-T
- Republish on Medium — ranks own URL, cannibalizes canonical, SKIPPED
- Short posts <1500w on blog — fails cluster depth + chunk-friendly structure

N - STRATEGIC DECISIONS LOG — choices we have made (2026-Q2)

Brand & identity

- Flagship = GraphRAG / Agentic KG (June 2026)
- Location = USA across platforms
- Mutassir personal + Cognitum org (light)
- Engineering-precision - light - blue mono-anchor design
- Voice = first-person engineer - no hype

Identity = the moat

Founder voice cannot be outsourced

Channels & platforms

- X / Twitter — opted OUT (platform direction)
- Medium — SKIPPED (cannibalizes canonical)
- LinkedIn mirror — PAUSED (Comm. Mgmt. API)
- Hashnode — DORMANT (paywall)
- Tier-2 news — HOLD until 20+ posts

Selective, not exhaustive

Investment matched to ROI window

Monetization & ops

- No monetization on site (RRM dormant)
- Consulting funnel is the only goal
- Skip proprietary token rotation
- Daily news cadence target (5+ /wk)
- Quarterly evergreen-blog refresh

Focus on lead quality, not ads

Brand-first, monetization later

Audience-first principles

- Newsletter is the only fully owned asset — top priority growth
- LinkedIn personal > LinkedIn corporate (engagement rates)
- Quality > volume — 1 deep post > 5 shallow
- Honesty over polish (Founder voice rule)
- Brand stays optimized (don't soften for everyone)

Audience built on substance, not reach hacks

Long-term thinking, not quarterly metrics

O - 90-DAY STRATEGY ROADMAP — phased synthesis of all the above

Days 1-30 - Foundation

- Hit 5wk news cadence → 32 posts (Tier-2 threshold)
- Publish GraphRAG pillar (currently in draft)
- LinkedIn personal: 2 posts/wk each (Mutassir + AI)

60+ tech-news posts (from 2) - 7 new posts

Days 31-60 - Acceleration

- Continue news cadence - add 1 blog/wk
- Build 2 lead magnets (GraphRAG benchmark, AI implementation)

END-OF-90-DAY TARGET STATE

100+ posts, 100+ backlinks, 100+ mentions, 100+ qualified consulting inbound, +35% AI mentions from net

Days 61-90 - Compounding + audit

- Continue cadence - add lead magnet 3 (State of Production AI Eng report)
- First podcast guest appearance (Latent Space or Pragmatic Engineer)
- 100+ backlinks, 100+ mentions, 100+ qualified consulting inbound

Cognitum AI - Research Knowledge Atlas - 2026-Q2 - sources: production CEO paper - openai/llm-ecosystem - google search central - no-research - Anthropic / openai/docs - openai/llm-ecosystem - llm-ecosystem - llm-ecosystem

The research knowledge atlas: the full body of SEO / GEO / AEO / E-E-A-T research the stack was built on, in one map.