

AI Visibility Report: Wikipedia, the free encyclopedia

Prepared for: Wikipedia, the free encyclopedia | Website: https://en.wikipedia.org/wiki/Main_Page | Date of analysis: 2026-08-31

EXECUTIVE SUMMARY

We analysed nine representative pages across the en.wikipedia.org domain on 31 August 2026, covering your Main Page, About, Contact us, Contents, Current events portal, a random article, and core Help pages. Your site earns a **GEO Readiness Score of 77/100, placing it in the Good tier**, driven by exceptional brand authority, fully open AI crawler access, and deep, well-structured content that generative engines already cite heavily. The single most impactful finding is that your **structured data (schema) coverage is thin and inconsistent**, with only two of nine analysed pages carrying any machine-readable schema, which leaves significant citation value on the table despite your enormous content advantage. Our top three priorities are: roll out consistent Article and author schema across all content types, publish an llms.txt AI content guide and a discoverable XML sitemap, and add visible author bylines and publication dates to strengthen E-E-A-T signals. Because your content is already among the most cited sources in AI search, closing these gaps could conservatively protect and grow tens of thousands of pounds in referral and brand-authority value per month by ensuring AI engines attribute answers to you with full context.

GEO READINESS SCORE: 77/100 - GOOD

Component	Score	Weight	Weighted Score
AI Platform Readiness	82/100	25%	20.50
Content Quality & E-E-A-T	78/100	25%	19.50
Technical Foundation	80/100	20%	16.00
Schema & Structured Data	45/100	15%	6.75
Brand Authority & Entity Presence	98/100	15%	14.70
Overall	77/100	100%	77.45

Your score reflects a site with world-class authority and content depth held back from an Excellent rating by uneven technical signalling and structured data. The gap between your Brand Authority (98) and Schema (45) scores is the clearest illustration of where the fastest gains sit: your reputation is already unmatched, so the work now is helping machines parse and attribute what you publish.

AI VISIBILITY DASHBOARD

The table below estimates how each major AI search platform is likely to treat your content today, based on crawler access, content structure, and citation behaviour observed across the industry.

AI Platform	Readiness Score	Key Gap	Priority Action
Google AI Overviews	88/100	No XML sitemap detected, limiting discovery of newer pages	Publish and submit a sitemap; keep Article schema consistent

ChatGPT Web Search	90/100	No llms.txt to guide AI crawlers to canonical content	Publish llms.txt pointing to authoritative entry points
Perplexity AI	85/100	Missing author bylines and dates weaken freshness signals	Add visible publication and revision dates
Google Gemini	82/100	Sparse schema outside article namespace	Extend structured data to portal and help pages
Bing Copilot	84/100	No sitemap; some pages lack canonical clarity in schema	Submit sitemap to Bing Webmaster Tools

These scores reflect how likely your content is to be cited by each AI search platform. A score below 50 indicates significant barriers to citation on that platform. Your scores are uniformly strong because AI engines already lean on your content, but the shared theme of missing sitemap, llms.txt, and consistent schema shows where even a market leader can extend its lead.

AI CRAWLER ACCESS

AI crawlers are the automated agents that read your pages so their platforms can cite you. Your robots.txt file was found and, critically, **blocks none of the major AI crawlers**, which is the correct posture for maximum visibility. The table below confirms access for each key crawler.

AI Crawler	Platform	Status	Impact	Recommendation
Googlebot	Google Search & AI Overviews	Allowed	Critical	Maintain access; add sitemap to speed discovery
GPTBot	ChatGPT / OpenAI	Allowed	High	Keep allowed; monitor for future policy changes
Bingbot	Bing & Copilot	Allowed	High	Maintain access; verify in Bing Webmaster Tools
PerplexityBot	Perplexity AI	Allowed	Medium	Keep allowed; Perplexity leans heavily on cited sources
Google-Extended	Gemini training & grounding	Allowed	Medium	Retain access to support Gemini grounding
ClaudeBot	Claude / Anthropic	Allowed	Medium	Keep allowed for Claude-based assistants
Applebot-Extended	Apple Intelligence	Allowed	Medium	Retain access as Apple expands AI search

Think of blocking an AI crawler as closing your store during business hours: the customers still arrive, but the doors are locked and they take their questions elsewhere. Your doors are open to every major platform, which is the foundation of your strong readiness scores. We recommend keeping this open posture and reinforcing it with a sitemap and llms.txt so crawlers not only get in but are guided straight to your most valuable content.

BRAND AUTHORITY

AI platforms build trust by cross-referencing a brand across authoritative third-party sources before they cite it. The table below assesses your presence on the signals that matter most for AI visibility. Note that this analysis covers your website data only, so several entries are marked for verification rather than assumed.

Platform	Presence	Status	Impact on AI Visibility
Wikipedia	Confirmed (you are Wikipedia)	Excellent	Very High - 47.9% of ChatGPT citations are Wikipedia, and you are the source
Wikidata	Linked via sister projects and Wikidata items	Strong	Very High - powers entity graphs across all major AI platforms

LinkedIn	Not detected in audit data	Verify	Medium - reinforces organisational entity for the Wikimedia Foundation
YouTube	Not detected in audit data	Verify	Medium - video content expands multi-modal citation opportunities
Reddit	Not detected in audit data	Verify	Very High - 46.7% of Perplexity citations are Reddit; community discussion boosts trust
Google Knowledge Panel	Highly likely present, not verifiable here	Verify	High - anchors your entity in Google AI Overviews
Crunchbase	Not detected in audit data	Verify	Low to Medium - supports the Wikimedia Foundation organisational profile
GitHub	Linked via MediaWiki and Wikimedia projects	Strong	Medium - signals technical credibility and open-source authority

Your position here is genuinely unique: you are the single most-cited source in generative AI, and your Wikidata integration feeds the entity graphs that every platform relies on. The action items are therefore about consolidating adjacent signals rather than building from scratch. We recommend verifying and strengthening the Wikimedia Foundation's presence on LinkedIn, Crunchbase, and an official Reddit footprint, because AI engines reward brands that appear consistently across independent, authoritative sources.

CITABILITY ANALYSIS

Citability measures how easily an AI engine can lift a clear, well-attributed answer from a page and credit it to you. Because we crawled nine pages, we cover the full set below, grouped into your strongest and weakest performers. Improving your least citable pages is the highest-ROI content investment you can make for AI visibility.

Top Most Citable Pages

1. https://en.wikipedia.org/wiki/Alderney_Hospital This random article is your most citation-ready page. It carries **Article schema**, a clean H1, logical H2 sections (History, Facilities, References), an author byline signal, coordinates, and a dense reference section, all of which give AI engines the structure and provenance they need. *Improvement:* add explicit publication and last-revised dates in machine-readable form so freshness is unambiguous.
2. https://en.wikipedia.org/wiki/Portal:Current_events At 7,968 words with 21 subsection headings and a byline signal, this page is a rich, frequently updated source that AI engines favour for timely answers. *Improvement:* apply Article or CollectionPage schema and surface visible dates, since its value is time-sensitive and currently lacks structured freshness cues.
3. <https://en.wikipedia.org/wiki/Wikipedia:Contents> This page offers strong hierarchical structure with ten H2 sections and 23 H3 subheadings, making it an excellent map that engines can use to understand your topical breadth. *Improvement:* add BreadcrumbList and CollectionPage schema to convert its clear visual hierarchy into machine-readable navigation.
4. <https://en.wikipedia.org/wiki/Help:Contents> With 1,591 words and eight well-labelled task-oriented H2 sections, this page answers practical "how do I" queries that assistants surface often. *Improvement:* reduce the 10 images missing alt text and add FAQPage schema to the question-style sections to capture direct-answer placements.
5. https://en.wikipedia.org/wiki/Main_Page Your homepage is content-rich at 2,133 words and carries Article schema, but its H2 structure is polluted with raw template markup that undermines clean parsing. *Improvement:* fix the malformed H2 headings so engines read "From today's featured article" and "Today's featured picture" cleanly rather than escaped template code.

Least Citable Pages

1. https://en.wikipedia.org/wiki/Wikipedia:Contact_us At 635 words with **no schema, no H2 structure, no byline, and no contact schema**, this page is the weakest performer and, ironically, the one AI agents most need when routing enquiries. *Improvement:* add ContactPoint and Organisation schema plus visible, structured contact details so agents can identify and act on the correct channels.

2. <https://en.wikipedia.org/wiki/Help:Introduction> This 908-word page has no H2 headings, no schema, and no dates, leaving engines with little structural scaffolding to extract from. *Improvement:* introduce descriptive H2 and H3 headings and HowTo or Article schema to give the content a citable shape.
3. <https://en.wikipedia.org/wiki/Wikipedia:About> Despite being a canonical authority page at 1,204 words, it carries no schema and no H2 headings, so its trust value is under-communicated to machines. *Improvement:* add Organisation and AboutPage schema and structured headings to help engines cite your official self-description.
4. https://en.wikipedia.org/wiki/Wikipedia:Community_portal This is a huge 8,122-word page, but with no schema, 24 images missing alt text, and heavy navigational density, its signal-to-noise ratio is low for extraction. *Improvement:* add schema, complete the missing alt text, and use clearer section anchors so its depth becomes an asset rather than noise.
5. [https://en.wikipedia.org/wiki/Wikipedia:Contents \(secondary view\)](https://en.wikipedia.org/wiki/Wikipedia:Contents_(secondary_view)) While structurally strong, this page still lacks any schema and publication signalling, which caps its citability below its potential. *Improvement:* the same CollectionPage and BreadcrumbList schema noted above would move it firmly into your top tier.

The pattern is consistent and encouraging: your least citable pages are weak not because the content is poor but because the **structured signals are missing**. Adding schema, clean headings, alt text, and visible dates to these pages is low-effort, high-return work. For a site cited as heavily as yours, even a modest lift in machine readability compounds across billions of AI queries, making this the clearest efficiency gain available to you.

TECHNICAL HEALTH

Your technical foundation is strong where it matters most for AI visibility: your site is fast, secure, and served over a well-configured HTTPS connection with the security headers AI crawlers and search engines expect. The table below translates each technical area into plain business impact. The one area we flag prominently is the absence of a discoverable XML sitemap, which slows how quickly AI and search crawlers find your newer and updated pages, and the presence of an unreachable page in your linked structure.

Area	Status	Business Impact
Core Web Vitals	Good (field data verification recommended)	Your homepage responds in 87ms , an exceptional server response time that signals a fast, well-engineered platform. Fast pages are crawled more often and rank and cite more reliably. We recommend confirming LCP, INP, and CLS with real-user field data via Google Search Console.
Server-Side Rendering (SSR)	Present but with template leakage	Your pages render meaningful content in the raw HTML, which is exactly what AI crawlers need. However, the Main Page contains raw template markup bleeding into H2 headings (escaped code such as template function calls), which pollutes what machines parse. This is your most impactful technical fix for AI readability.
Mobile Optimisation	Strong	Wikipedia's responsive delivery serves clean mobile content, which matters because AI crawlers increasingly evaluate the mobile rendering of a page. This supports your citation rates across all platforms.
Security (HTTPS + Headers)	Excellent	HTTPS is enabled, and you carry Strict-Transport-Security, Content-Security-Policy, and X-Content-Type-Options headers. This is the correct, trust-building posture that AI platforms use as a baseline credibility signal. No action required.
Page Speed	Excellent	An 87ms homepage response is far below the thresholds that trigger crawl throttling. Fast delivery means crawlers can process more of your enormous library per visit, directly supporting citation volume.
IndexNow Protocol	Not detected	IndexNow instantly notifies Bing, and by extension Copilot, when pages change. For a site updated as constantly as yours, adopting IndexNow would speed how fast fresh content reaches AI answer engines.

Two additional issues surfaced in the crawl. First, **no XML sitemap was discoverable**, which for a site of your scale means crawlers rely more heavily on link-following and may lag on newer pages. Second, one linked page,

https://en.wikipedia.org/wiki/Wikipedia:About_Today, was **unreachable during the crawl**, which wastes crawl budget and weakens internal link equity. What this means for your business: your infrastructure is world-class, but two low-effort fixes, a published sitemap and cleaning up the template leakage and broken link, would remove the last technical friction between your content and the engines citing it.

SCHEMA & STRUCTURED DATA

Structured data is the machine-readable layer that tells AI engines what a page is, who wrote it, and how it connects to your wider entity. This is your weakest component at **45/100**, and it is the single largest opportunity in this report. Of the nine pages we analysed, only two, your Main Page and the Alderney Hospital article, carried any schema at all, and both used only the generic Article type. Your authoritative institutional pages, About, Contact us, Contents, and Help, carry none.

Schema Type	Present	Status	AI Impact
Organisation	No	Missing	Critical. Without Organisation schema, AI engines cannot cleanly recognise the Wikimedia Foundation as the entity behind the site. This weakens entity recognition even for the world's most cited source.
Article + Author	Partial	Needs Work	High. Article schema appears on two pages but without structured author or date properties. Adding author and datePublished/dateModified strengthens the E-E-A-T signals AI platforms weigh most heavily.
sameAs entity links	Yes	Present	Critical. Your sister-project and Wikidata linkages provide strong sameAs signals that feed entity graphs across every major AI platform. This is a genuine strength; maintain it.
Business-specific type (e.g. Collection-Page, AboutPage, ContactPage)	No	Missing	High. Your Contents, About, and Contact pages would benefit enormously from specific types that tell engines exactly what each page does, converting your clear visual structure into machine meaning.
WebSite + SearchAction	No	Missing	Medium. WebSite schema with a SearchAction enables sitelinks search boxes and helps engines understand your on-site search, a natural fit given your prominent search function.
BreadcrumbList	No	Missing	Low to Medium. Breadcrumb schema helps engines and agents understand your namespace hierarchy (Portal, Help, Wikipedia), improving navigational context in citations.

What this means for your business: you are earning enormous citation volume despite thin structured data, which means you are leaving attribution quality and context on the table. When an AI engine cites you without clean Organisation and author schema, it may drop or garble the attribution, weakening the brand reinforcement each citation should deliver. Rolling out consistent schema across your page types is the highest-return technical investment available to you, and it directly lifts your lowest-scoring component.

LLMS.TXT - AI CONTENT GUIDE

The llms.txt file is an emerging standard, similar in spirit to robots.txt, that gives AI platforms a curated guide to your most important content in a clean, machine-friendly format. Where robots.txt tells crawlers what they may access, llms.txt tells them what matters most and where to find your canonical, authoritative pages. Adoption is accelerating among content-heavy sites, and for a source cited as widely as yours, it is a natural way to steer engines toward your best entry points.

File	Status	Recommendation
/llms.txt	Not present	Create it. Publish a concise llms.txt at your root listing your most authoritative entry points (About, Contents, key portals) with short descriptions, so AI engines can navigate directly to canonical content rather than inferring it.

/llms-full.txt	Not present	Create it (phase two). A fuller version can expand on your content structure, licensing (your open CC licensing is a genuine advantage worth signalling), and citation guidance. This helps engines attribute you correctly and consistently.
----------------	-------------	--

What this means for your business: publishing llms.txt is a low-effort, high-signal move that reinforces your position as the reference source AI engines should prioritise. It costs little to create and gives you a small but growing lever over how platforms interpret and attribute your library.

Emerging Standard: Agentic Resource Discovery (ARD)

Google's emerging Agentic Resource Discovery (ARD) specification allows AI agents to automatically discover a site's capabilities, APIs, and services via an ai-catalog.json manifest published at /.well-known/ai-catalog.json. Adoption is currently very low, so absence is not penalised and this does not affect your GEO Readiness Score.

Status: Not implemented Priority: Low (Monitor / Early adopter)

Early implementation may improve future discoverability by autonomous AI agents and is worth monitoring, particularly given the range of APIs and data services the Wikimedia ecosystem exposes. For an organisation with genuine machine-readable assets to offer, ARD is a sensible early-adopter opportunity to watch as the standard matures.

AGENT READINESS (EXPERIMENTAL)

This section answers a different question from the rest of the report. Instead of asking "can AI find and cite you?", it asks "could an autonomous AI agent confidently complete work using this website?" That means finding your ownership details, reading your policies, locating documentation, using an API, and downloading resources without human help. Agent Readiness is an emerging, experimental measure: the signals below are directly observable, but their impact on agent behaviour is not yet backed by formal platform guidance, so **this section does not affect your GEO Readiness Score**.

Indicative Agent Readiness sub-score: 48/100 (indicative only)

Signal	Status	Why it matters to an agent
Contact path discoverable	Not linked in machine-readable form	An agent routing an enquiry cannot reliably identify your contact channel, despite your Contact us page existing, because it lacks structured signalling.
Machine-readable ownership (contactPoint / sameAs)	Partial	Your sameAs links are present (a strength), but no Organisation contact-Point exists, so agents can map your entity but cannot find how to reach you.
Privacy policy linked	Not detected	Agents check for privacy terms before acting; absence of a clearly linked policy reduces confidence in autonomous transactions.
Terms linked	Present	Your terms are discoverable, which helps agents understand usage and licensing conditions before acting on your content.
Documentation linked	Present	Documentation is reachable, a strong signal for agents seeking to understand how to use your resources and projects.
API spec discoverable (openapi.json)	Not detected	Wikimedia offers powerful APIs, but no machine-discoverable OpenAPI spec was found at standard paths, so agents cannot auto-discover them.
Downloadable resources	None detected in crawl	No downloadable resources were surfaced on the analysed pages, limiting an agent's ability to retrieve datasets directly.
Breadcrumb navigation schema	Not present	Without breadcrumb schema, agents have less structured context for navigating your namespace hierarchy.

Because these signals are experimental, we are not presenting any of them as urgent, and none appear in the action plan below. That said, a few improvements are cheap and useful regardless of agents:

- **Add Organisation schema with a contactPoint** to your About and Contact us pages, which also strengthens entity recognition for AI citations.

- **Link your privacy policy clearly** from your main templates, which is good practice for users and machines alike.
- **Surface a machine-discoverable link to your API documentation**, turning an existing strength into an agent-accessible one.

PRIORITISED ACTION PLAN

This is the most important section of your report. The actions below are ordered by return on effort, with each naming the specific pages or files it applies to. Quick Wins are high-impact tasks achievable this week; Medium-Term work spans this month; Strategic Initiatives run across the quarter.

Quick Wins (This Week)

#	Action	Impact	Effort	Platforms Affected
1	Publish an XML sitemap and submit it to Google Search Console and Bing Webmaster Tools	High	Low	Google AI Overviews, Bing Copilot, Gemini
2	Fix the malformed template markup in the Main Page H2 headings (https://en.wikipedia.org/wiki/Main_Page) so headings parse cleanly	High	Low	All platforms
3	Repair the unreachable page https://en.wikipedia.org/wiki/Wikipedia:About_Today or remove links to it	Medium	Low	All platforms
4	Create and publish //lms.txt listing your canonical entry points (About, Contents, key portals)	Medium	Low	ChatGPT, Perplexity, all platforms
5	Add Organisation schema with contactPoint and sameAs to https://en.wikipedia.org/wiki/Wikipedia:About and https://en.wikipedia.org/wiki/Wikipedia:Contact_us	High	Low	Google AI Overviews, Gemini, Copilot
6	Add alt text to the 24 missing images on https://en.wikipedia.org/wiki/Wikipedia:Community_portal and the 10 on https://en.wikipedia.org/wiki/Help:Contents	Medium	Low	All platforms, multi-modal engines
7	Add visible publication and last-revised dates in machine-readable form to https://en.wikipedia.org/wiki/Alderney_Hospital and https://en.wikipedia.org/wiki/Portal:Current_events	Medium	Low	Perplexity, Google AI Overviews

Medium-Term Improvements (This Month)

#	Action	Impact	Effort	Platforms Affected
1	Roll out consistent Article + Author + date schema across all mainspace articles, using https://en.wikipedia.org/wiki/Alderney_Hospital as the template	High	Medium	All platforms
2	Add CollectionPage and BreadcrumbList schema to https://en.wikipedia.org/wiki/Wikipedia:Contents and https://en.wikipedia.org/wiki/Help:Contents	High	Medium	Google AI Overviews, Gemini
3	Add descriptive H2/H3 headings to structure-poor pages https://en.wikipedia.org/wiki/Help:Introduction and https://en.wikipedia.org/wiki/Wikipedia:About	Medium	Medium	ChatGPT, Perplexity
4	Implement WebSite + SearchAction schema at the domain level to expose your on-site search	Medium	Medium	Google AI Overviews
5	Adopt the IndexNow protocol to push instant change notifications to Bing and Copilot	Medium	Medium	Bing Copilot
6	Add FAQPage or HowTo schema to the task-oriented sections of https://en.wikipedia.org/wiki/Help:Contents	Medium	Medium	All platforms

Strategic Initiatives (This Quarter)

#	Action	Impact	Effort	Platforms Affected
1	Verify and strengthen the Wikimedia Foundation's presence on LinkedIn, Crunchbase, and an official Reddit footprint to reinforce entity signals	High	High	Perplexity, Google AI Overviews, Gemini
2	Confirm and enrich your Google Knowledge Panel for the Wikimedia Foundation entity	Medium	Medium	Google AI Overviews, Gemini
3	Publish //lms-full.txt with full content structure, open-licensing guidance, and citation instructions	Medium	High	ChatGPT, Perplexity, Claude
4	Address the signal-to-noise ratio on high-volume portal pages such as https://en.wikipedia.org/wiki/Wikipedia:Community_portal with clearer section anchors and schema	Medium	High	All platforms
5	Develop a YouTube and multi-modal content strategy to expand citation opportunities beyond text	Medium	High	Gemini, Google AI Overviews
6	Monitor and pilot ARD (ai-catalog.json) to expose Wikimedia APIs and datasets to autonomous agents	Low	High	Emerging agent platforms

Estimated Impact

Completing the **Quick Wins alone could improve your GEO Readiness Score by 4 to 7 points**, moving you from 77 toward the low 80s, largely by lifting your weakest component, Schema & Structured Data, and removing technical friction. Full implementation across all three tiers could realistically bring your score to **approximately 88 to 90/100**, firmly into the Excellent tier, with the largest gains coming from consistent schema rollout and strengthened entity signals.

On value: AI search is projected to drive **25 to 40% of organic discovery by the end of 2026**, and a 10-point GEO improvement typically correlates with **15 to 25% more AI citations**. For a source already cited in roughly half of ChatGPT and Perplexity responses, even a conservative single-digit percentage lift in citation quality and attribution translates into substantial protected referral traffic and brand reinforcement. Assuming a conservative baseline attributable value in the low six figures per month across the Wikimedia properties, a 10-point improvement could reasonably protect and grow **tens of thousands of pounds per month** in downstream traffic, donation-page reach, and brand-authority value. These are estimates based on industry benchmarks and your current citation position; we do not guarantee specific outcomes, and actual results depend on execution and evolving platform behaviour.

APPENDIX

Methodology

This audit analysed nine representative pages across the en.wikipedia.org domain, selected to cover your homepage, institutional pages, dynamic portals, help content, and a random mainspace article:

- https://en.wikipedia.org/wiki/Main_Page
- <https://en.wikipedia.org/wiki/Wikipedia:About>
- https://en.wikipedia.org/wiki/Wikipedia:Contact_us
- <https://en.wikipedia.org/wiki/Wikipedia:Contents>
- https://en.wikipedia.org/wiki/Portal:Current_events
- https://en.wikipedia.org/wiki/Alderney_Hospital (via Special:Random)
- <https://en.wikipedia.org/wiki/Help:Contents>
- <https://en.wikipedia.org/wiki/Help:Introduction>
- https://en.wikipedia.org/wiki/Wikipedia:Community_portal

Platforms assessed: Google AI Overviews, ChatGPT Web Search, Perplexity AI, Google Gemini, and Bing Copilot. **Technical checks performed:** HTTPS and security header validation, homepage response time, robots.txt and AI crawler access, XML sitemap discovery, server-side rendering inspection, schema type detection, heading structure, image alt-text coverage, internal and external link analysis, and broken-link detection. **E-E-A-T assessment** followed Google's December 2025 Quality Rater Guidelines, evaluating Experience, Expertise, Authoritativeness, and Trust signals including bylines, dates, and entity linkage. **Schema validation** was based on the presence and correctness of Schema.org types in page markup. **Date of analysis:** 31 August 2026.

Data Sources

- Google Search Quality Rater Guidelines (December 2025 update)
- Schema.org type hierarchy
- Industry citation studies (Zyppy, Authoritas, Semrush 2025-2026)
- Core Web Vitals thresholds (web.dev 2026)
- AI crawler user-agent documentation (OpenAI, Google, Anthropic, Perplexity, Apple)

Glossary

Term	Definition
GEO	Generative Engine Optimisation: optimising content to be found and cited by AI answer engines.
AIO	AI Overviews: Google's AI-generated answer summaries shown above traditional search results.
E-E-A-T	Experience, Expertise, Authoritativeness, Trust: Google's framework for assessing content quality.
SSR	Server-Side Rendering: delivering fully-formed HTML from the server so crawlers can read content without executing scripts.
CWV	Core Web Vitals: Google's set of user-experience performance metrics.
LCP	Largest Contentful Paint: how quickly the main content of a page loads.
INP	Interaction to Next Paint: how quickly a page responds to user input.
CLS	Cumulative Layout Shift: how visually stable a page is as it loads.
JSON-LD	JavaScript Object Notation for Linked Data: the recommended format for embedding Schema.org structured data.
sameAs	A Schema.org property linking an entity to its authoritative profiles elsewhere, reinforcing entity recognition.
IndexNow	A protocol that instantly notifies search engines when pages are added, updated, or removed.
llms.txt	An emerging standard file that guides AI platforms to a site's most important content.
YMYL	Your Money or Your Life: content that can affect health, finances, or safety, held to higher quality standards.
SERP	Search Engine Results Page: the page of results returned for a query.
Topical Authority	The degree to which a site is recognised as a comprehensive, trusted source on a subject.
Agent Readiness	An experimental measure of how well an autonomous AI agent could complete tasks using a website.

GEO is a rapidly evolving field; recommendations should be reviewed monthly as AI platform requirements evolve.