



VOLUME ONE · 2026 INDUSTRY REPORT

Search, *Reframed.*

A field manual for marketers, founders and agencies operating in a search environment that no longer ends with a blue link.

SKOMA DIGITAL
SKOMADIGITAL.COM.AU

AUTHORED BY
SAMI KHAN, PRINCIPAL
ADVISOR

PUBLISHED
MAY 2026

A note before you begin.

This report was written for the marketers, founders, and agency leaders who have been told, repeatedly and with increasing intensity, that search is broken, that SEO is dead, that everything they have built is about to be displaced by something called GEO. We do not believe any of that, and we suspect you do not believe most of it either.

What is true is that the surface of search has changed. The interface a person sees when they ask Google a question in May 2026 looks very little like the ten blue links that defined the previous twenty years. AI Overviews, AI Mode, ChatGPT Search, Perplexity, Claude, Gemini, Bing Copilot, and a fast-multiplying ecosystem of AI assistants now sit between the user and the open web. For some queries they answer the question entirely. For others they shortlist sources for the user to examine.

What has not changed is the underlying economics. Brands that command authority, expertise, structured information, and trust are still the brands that get cited. The signals look familiar to anyone who has worked through a decade of search: relevance, authority, freshness, structure, evidence. What has shifted is how those signals are read, by whom, and how the resulting visibility is measured.

Most of the AI search commentary you will encounter in 2026 falls into two camps. The first sells fear, telling you the open web is collapsing and you must rebuild your stack tomorrow. The second sells a tactical sleight of hand, suggesting that a few schema tags and an l1ms.txt file will buy you visibility inside ChatGPT. Neither view is honest with the evidence.

Our intent here is to offer a third position. To present the verified data from Google, Anthropic, OpenAI, Microsoft, Pew Research, the Reuters Institute, and the academic record with the level of care a major business decision deserves. To distinguish what has structurally changed from what has merely been rebadged. To give you the language, the frameworks, and the action map you need to lead a credible response, not a reactive one. You will find places in this report where we disagree with vendor talking points, including some we would commercially benefit from agreeing with. Where the data is contested we have said so. Where the trend is real but the scale is small, we have said that too. The aim has been to give you something you can hand to your board and to your team without needing to caveat it.

The transition is not finished. Anyone who tells you it is, or that they know how it ends, is selling something. What follows is our best read of where the discipline now stands, and what the next two to five years are likely to require. Thank you for the time you are about to spend with this work. If any part of it changes how you think, that has been the entire point.

Sami Khan

PRINCIPAL ADVISOR & FOUNDER · SKOMA DIGITAL

CONTENTS

What is inside.

FRONT MATTER

•	From the Founder	02
	An honest framing of what has changed	
•	Executive Summary	04
	The ten findings that matter most	

PART ONE. THE ARCHITECTURE SHIFT

01	The State of Search in 2026	07
	Where the surface ends and the engine begins	
02	From SEO to GEO	12
	The discipline, reframed	
03	How AI Systems Discover Information	16
	Retrieval, embeddings, knowledge graphs	
04	Understanding AI Retrieval and Citations	20
	What gets quoted, and why	

PART TWO. VISIBILITY, AUTHORITY, OPERATIONS

05	AI Search Platforms Compared	24
	Five engines, five citation logics	
06	The New Visibility Framework	28
	Five layers of presence	
07	Entity Authority and Trust	32
	Why some brands repeat in every answer	
08	Content Systems for AI Discovery	36
	What the Princeton study actually proved	
09	Technical SEO in AI Search	42
	Crawlers, schema, the bot policy you have not audited	

PART THREE. STRATEGY AND OUTLOOK

10	Human Expertise versus AI Scale	46
	The quality bar, raised	
11	What Businesses Are Getting Wrong	49
	Eleven mistakes the data already records	
12	AI Workflow Systems for Modern Marketers	52
	Editorial governance for the AI-assisted team	
13	Future of Search and Discovery	56
	Agents, memory, and the next two to five years	
14	Practical Action Framework	60
	A six-step plan for the next ninety days	

REFERENCE MATERIAL

•	Strategic Checklist	63
•	Key Takeaways	66
•	Full Source Library	67

What this report concludes, in ten findings.

Search did not break in the AI era. It fragmented, and the criteria for visibility consolidated around a smaller set of authority signals. The brands willing to engineer for those signals are pulling ahead. The brands waiting for clarity are losing ground they will not recover quickly.

The ten findings that follow are the conclusions our research, and the verified record from Google, OpenAI, Anthropic, Microsoft, Pew Research, the Reuters Institute, BrightEdge, Ahrefs, Semrush, Similarweb and the peer-reviewed academic literature, support today. Each is expanded across the chapters that follow.

1. AI search did not replace traditional search. It is layered on top of it.

Google still processes roughly five trillion searches per year and its market share remains close to ninety per cent. AI Overviews and AI Mode are now part of Google Search rather than alternatives to it. ChatGPT, Perplexity, Gemini, Claude and Bing Copilot have grown materially in 2024 and 2025 but they have done so alongside, not instead of, conventional search. The competitive picture is additive, not substitutive.

2. The click economics have changed, but the brand economics have not.

Pew Research, in a sample of 68,879 searches collected in March 2025, found that when a Google search result page contained an AI summary, users clicked through to a traditional link 8 per cent of the time, compared with 15 per cent without one. Only 1 per cent clicked an AI Overview citation. Sessions ending without any further action rose from 16 to 26 per cent.

Visibility, citation share and brand mention rate are now leading indicators that conventional analytics dashboards under-measure. The brands that adapt their measurement framework first will be the ones that can defend their budgets when click counts continue to compress.

3. Generative Engine Optimisation is authority engineering, not formatting tricks.

The often-quoted Princeton, Georgia Tech and IIT Delhi study published at ACM SIGKDD 2024 tested nine content modifications across ten thousand queries on five generative engines. Adding citations, statistics, and quotations from authoritative sources improved AI citation visibility by between thirty and forty per cent. Keyword stuffing reduced visibility by ten per cent. The pattern that emerges is not stylistic. It is that LLMs reward content that signals evidence and authority.

+40%

Average lift in AI citation visibility for content that includes statistics, citations and quotations from authoritative sources, across ten thousand queries on five generative engines.

SOURCE. AGGARWAL, MURAHARI, RAJPUROHIT, KALYAN, NARASIMHAN, DESHPANDE. GEO: GENERATIVE ENGINE OPTIMIZATION, ACM SIGKDD KDD 2024.

4. Each AI engine cites differently. Multi-engine measurement is now the baseline.

Profound, in an analysis of 680 million citations between August 2024 and June 2025, found only 11 per cent overlap in the top sources cited by ChatGPT and Perplexity. Wikipedia accounts for roughly 7.8 per cent of ChatGPT citations. Reddit dominates Perplexity and Google AI Overviews in many categories. YouTube takes 29.5 per cent of Google AIO citations across studies surveyed by Lantern in February 2026. Optimising for one engine is not optimising for all.

5. The same E-E-A-T signals do double duty in organic and AI surfaces.

BrightEdge data published in September 2025 found that 54.5 per cent of Google AI Overview citations now come from pages already ranking in the organic top ten, up from 32 per cent twelve months earlier. The investments that build organic authority compound across both interfaces. The split that some agencies are drawing between SEO and GEO is largely artificial.

6. The Australian rollout lag is a window. It is closing.

AI Overviews launched in the United States in May 2024 and in Australia at the end of October 2024, a five-month gap. AI Mode launched in the United States in May 2025 and in Australia in October 2025. Australian brands have, for the moment, the rare advantage of seeing a more advanced market move first. That observational lead has shrunk to roughly five months and is expected to disappear over the next eighteen.

7. Most brands are blocking the AI bots they should be allowing.

Cloudflare's 2025 crawler audit identified GPTBot, Meta-ExternalAgent, ClaudeBot and Bytespider as the four largest AI crawlers. Many brands, on legitimate copyright concerns, have blocked them at the WAF or robots.txt layer. The unintended consequence is that the search-and-citation bots that drive visibility inside ChatGPT, Claude and Perplexity, which are different bots from the training crawlers, have also been blocked. The number of brands that have inadvertently removed themselves from the most lucrative AI citation surfaces is, on the audits we have conducted, materially higher than the affected teams realise.

8. The AI content gold rush is over. The expert content era has begun.

Google's March 2024 core update, expanded in November 2024 and again in January 2025, re-defined scaled content abuse to apply regardless of whether content was produced by humans, AI, or a combination. The site reputation abuse policy demoted parasite SEO subdomains on major publishers including CNN, Forbes, the Wall Street Journal and Wired. Independent niche sites that had relied on programmatic content lost between sixty and ninety-five per cent

of their traffic in the affected period. Recovery, in the few documented cases, has only come for sites that can demonstrate genuine first-hand expertise.

9. Reddit, Wikipedia, YouTube and LinkedIn are the under-priced visibility surfaces.

Reddit's licensing agreement with Google, valued at approximately sixty million dollars per year, and a separate agreement with OpenAI of roughly seventy million dollars per year, have entrenched Reddit content inside the AI surfaces that matter most. The cheapest visibility increase available to most brands today is not paid media. It is credible, sustained presence on the platforms LLMs already trust.

10. Agentic browsing will reset the next layer of competition. Not yet, but soon.

OpenAI's Operator, integrated into ChatGPT agent in July 2025, scored 87 per cent on the Web-Voyager benchmark. Google's Project Mariner scored 83.5 per cent and is now embedded in the Gemini Enterprise Agent Platform announced at Cloud Next 2026. Anthropic's Computer Use scored 56 per cent. Reliability still trails human performance materially, but the direction is clear. Within two to five years, brand presence inside agent shortlists, on-device assistants and persistent AI memory systems will become a meaningful share of high-intent visibility.

The brands that win the next decade of search will not be the ones that move first. They will be the ones that build authority slowly enough to be defensible.

SKOMA EDITORIAL POSITION

What this report does not claim, and what it does

We do not claim search is dying. We do not claim AI traffic is about to overtake organic in volume. We do not claim that any one tool, llms.txt file, schema configuration or AI-content stack will solve the problem. The discipline of search is, as it has always been, the discipline of building durable signals of authority and serving real intent. The interface has changed. The work has not. The chapters that follow walk through how AI search engines retrieve information, how citation decisions are made, how brands can adapt their content systems and technical foundations, what is failing in the market today, and what the next two to five years are likely to require. A detailed action framework, a strategic checklist, a tearsheet of takeaways, and the full source library appear at the end.

01

The State of Search in 2026.

A measured account of where the surface ends, where the engine begins, and what has actually changed for the people who do this work for a living.

The surface changed. The engine did not.

Two interfaces now share the same retrieval index. The behaviour of the user, the architecture of the platform, and the economics of the click have all moved at different speeds. Untangling them is the precondition for any sensible strategy.

For most of the period from 2003 to 2023, search worked as a near-static idea. A user typed a query, an engine returned a list, the user clicked a link, the brand received a visit and could measure what happened next. Inside Google specifically, the underlying systems became radically more sophisticated over time, with PageRank giving way to BERT, MUM, RankBrain and dozens of other components, but the surface a user saw remained unchanged in its essential form. Ten blue links. A snippet. An ad rail.

What changed in 2024 was the surface, not the engine beneath it. Google AI Overviews launched in the United States on 14 May 2024 and reached Australia on 29 October 2024. AI Mode, a more conversational, multi-step variant, followed in the United States on 20 May 2025 and arrived in Australia on 8 October 2025. Both sit on top of the Google Search index. Both still depend on Googlebot crawling, on the Knowledge Graph, on the same E-E-A-T quality signals that govern conventional rankings. They are interfaces that consume the engine, not replacements for it.

The same architectural principle holds across the rest of the AI search ecosystem. ChatGPT Search uses Bing's index as its primary retrieval substrate. Perplexity has built proprietary indexing, but the underlying logic remains retrieval-augmented generation over the open web. Brave Leo uses the Brave Search index. Bing Copilot is a layered interface over Bing's grounded retrieval. None of these systems generates citations from the model's training data alone. They retrieve, synthesise, and cite.

What zero-click really means in 2025 and 2026

The most quoted statistic in AI search commentary is that some large fraction of searches now end without a click. The number ranges, in published studies, from 58.5 per cent (SparkToro and Datos, July 2024, US Google searches) to 69 per cent (Similarweb, 2025) to 83 per cent (Semrush, AI Overview-triggered queries, September 2025) to 93 per cent (Semrush, AI Mode queries, September 2025). The variance reflects methodology rather than disagreement. Different panels, different devices, different inclusion criteria, different definitions of a click.

The honest framing is this. Zero-click search has been growing for many years. The introduction of Featured Snippets, Knowledge Panels, People Also Ask, video carousels and shopping units inside the SERP all contributed long before AI surfaces existed. AI Overviews accelerated a trend that was already running. They did not create it.




FIGURE 1.1 Zero-click search has been a long arc, not a 2024 inflection.

Source. SparkToro and Datas, 2018 to 2024 series; Similarweb 2025. Methodologies differ year-on-year. Read as directional.

The Pew Research Centre study published in July 2025 is, in our view, the most carefully constructed data point on AI Overview impact available in the public record. Pew worked with a panel of nine hundred United States adults and observed 68,879 unique searches across March 2025. It found that on Google searches that included an AI summary, users clicked through to traditional links 8 per cent of the time, compared with 15 per cent on searches without one. Only 1 per cent clicked an AI Overview source link. The share of search sessions that ended browsing entirely rose from 16 to 26 per cent.




FIGURE 1.2 Click rates on Google searches with and without an AI summary.

Source. Pew Research Centre, July 2025. Sample of nine hundred US adults observed across 68,879 unique searches in March 2025.

Google has publicly disputed the framing of these numbers, arguing that AIO is broadening total search volume and that the clicks AIO does generate are of higher quality. Both can be true. The aggregate effect for any single brand depends on the distribution of its query mix and on whether its content is being cited inside the AIO when one appears.

The conversational shift, in the data

The Reuters Institute Digital News Report 2025, drawing on a survey of forty-eight markets, recorded weekly use of any AI tool moving from 18 per cent to 34 per cent year-on-year. Weekly use of ChatGPT specifically reached 22 per cent across the six markets where the question was asked. Among under-twenty-fives the number is materially higher. Reuters separately found that 9 per cent of global respondents now use AI chatbots for news, with the figure rising to between 18 and 27 per cent in the youngest age band.

OpenAI's own disclosures, supported by Sensor Tower's tracking, place ChatGPT at approximately 800 million weekly active users by October 2025. Google's Gemini grew faster than any other large-language assistant in the second half of 2025 and reached an estimated 450 million monthly active users. Perplexity and Claude remain smaller in absolute terms but are growing in their respective specialisations.




FIGURE 1.3 Weekly active users for the four largest AI assistants, January 2024 to October 2025.

Source. Vendor disclosures collated against Sensor Tower estimates. October 2025 figures are best public estimate; OpenAI confirmed 800 million weekly users.

Multimodal and visual

Less discussed but at least as important is the rise of multimodal search. Google reported at I/O 2025 that Lens now handles approximately twenty-five billion queries per month, up from twenty billion in October 2024. Search Live, the live-video and voice variant of Project Astra, is in active rollout. Pinterest Lens, Snapchat's visual search, and the camera-first behaviour of TikTok all point in the same direction. The text query, while not vanishing, is no longer the only entry point.

What this means for marketers

Three implications follow from the data above. We will return to each in subsequent chapters, but it is worth stating them clearly here.

The unit of competition has shifted from URL-rank to citation share. Where you appear inside an answer matters as much as whether you appear at the top of a list. Tracking systems that measure only ranking, not citation, are now incomplete.

The first move is measurement, not content. Many brands are already producing the right content. They have not yet adapted their dashboards to see what that content is achieving in AI surfaces. Bing Webmaster Tools added an AI Performance dashboard in February 2026. Google Search Console has added partial AI Overview impression data. Independent tooling from Profound, Otterly, Peec AI, AthenaHQ, Scrunch and BrightEdge AI Catalyst has matured rapidly. The first ninety days of any GEO programme should establish a baseline, not produce more content.

Authority is now a scarcer commodity than ever. When the surface compresses information into a synthesised answer, only a small set of sources is named. Being one of those sources requires investments that compound over years, not quarters. The brands that started building entity, schema, original research and expert authority programmes in 2022 and 2023 are now disproportionately visible. The brands that begin in 2026 will start to see results in 2027 and 2028. The work is real, and it cannot be sprinted.

The interface that meets the user is what changed. The discipline of building authority did not.

SKOMA EDITORIAL POSITION

02

From SEO to GEO.

What the discipline is being renamed, what is genuinely new, and where the renaming is more enthusiastic than accurate.

The discipline, reframed.

Generative Engine Optimisation is a useful term. It is not, however, a new discipline. It is the application of well understood search principles to a new set of interfaces, and the most useful first move is to be clear about what is genuinely changing.

The phrase Generative Engine Optimisation was introduced into the academic record by Aggarwal and colleagues at Princeton, Georgia Tech and IIT Delhi in November 2023, in a paper later presented at the ACM SIGKDD Conference on Knowledge Discovery and Data Mining in August 2024. Their definition is precise. GEO is the practice of optimising content so that it is more likely to be cited within the responses of generative engines, where a generative engine is any system that produces a synthesised answer rather than a list of links.

Bing was the first major search platform to formally recognise the term. In February 2026 Microsoft updated its Bing Webmaster Guidelines to include GEO by name and to expand its abuse definitions for AI-engineered language. Google has not yet adopted the acronym in its public guidance, although AI Mode and AI Overviews are now described in Search Central documentation in language that closely tracks what GEO practitioners would call optimisation.

What is structurally different

Three things are structurally different in the GEO surface, and three are functionally the same. We start with the differences.

The unit of victory is a citation, not a click. In conventional search, a brand wins by occupying a top organic or paid position and capturing a visit. In a generative engine, a brand wins by being cited inside the synthesised answer, and the visit is optional and often does not occur. This forces a re-think of measurement, attribution, and content design. A page that ranks third in Google but is the named source inside the AI Overview at position one is performing a different and arguably more valuable function.

Information is unbundled at the chunk level. Conventional retrieval ranks pages. AI retrieval, in production, almost always ranks chunks. A two-thousand-word page might be ingested as ten distinct three-hundred-token chunks, each independently scored against a query, and only one or two of those chunks may end up in the synthesised answer. The implication is that information density at the paragraph level matters more than at the page level. Sprawling, narrative pages that carry their main answer in the conclusion are at a disadvantage. Front-loaded, claim-rich pages are favoured.

Brand and entity recognition is the gating signal. A LLM does not browse the web. It retrieves a candidate set and then must decide which sources to attribute to which claims. That decision is materially influenced by whether the model recognises the source as authoritative within the topic. Brands that exist coherently inside the Knowledge Graph and Wikidata, that are linked from authoritative third-party sources, and that have a consistent semantic footprint, get cited disproportionately.

What is functionally the same

Three things are functionally the same as conventional SEO, even though the surface has changed.

Indexing is still the gate. A page that is not indexed cannot be retrieved, cannot be embedded, and cannot be cited. The technical SEO foundations of crawlability, sitemap hygiene, internal linking, indexability and canonical management are unchanged in their importance. The vendors selling GEO services that ignore these foundations are selling a partial product.

Authority is still cumulative. The same E-E-A-T signals that Google has been refining since 2014, and which Google explicitly extended in December 2022 with the addition of Experience as a fourth pillar, and again in September 2025 with explicit AI Overview evaluation criteria, govern AI citation behaviour. BrightEdge's longitudinal research, charted overleaf, makes the point quantitatively. Across the first sixteen months of AI Overviews in the United States, the share of AIO citations sourced from organic top-ten ranking pages rose from approximately 32 to 54.5 per cent.




FIGURE 2.1 The convergence of organic ranking and AI Overview citation, May 2024 to September 2025.

Source. BrightEdge Generative Parser, monthly tracking. Indicates that the AI surface is increasingly drawing from the same authority pool as conventional search.

Intent still wins. A page that solves the user's actual question, with appropriate evidence and clarity, is rewarded by both conventional and AI surfaces. The shape of the page may need to change. The fundamental contract between content and intent has not.

Where the renaming gets ahead of the evidence

Three claims about GEO are circulating in the market that we believe overstate the case.

The first is that conventional SEO is dead. It is not. Google's search volume continues to grow. Organic traffic remains, for almost every brand we work with, the largest single source of inbound discovery. AI traffic is growing fast in percentage terms but starts from a small base. According to Similarweb's 2025 data, AI referrals to the top one thousand sites totalled approximately 1.13 billion visits in June 2025, up 357 per cent year on year. That is meaningful, but it sits against an estimated 60 to 80 billion monthly visits to the same sites from conventional search.

The second claim is that adding an llms.txt file is a precondition for AI visibility. The llms.txt proposal, drafted by Jeremy Howard at Answer.AI in September 2024, suggests a convention by which a website can publish a Markdown index of its highest-value pages and a curated subset of its content for LLM ingestion. It is a useful idea. It is not, however, used by any of the major LLM platforms as a primary retrieval signal. Anthropic, OpenAI, Google and Microsoft have not endorsed it as a citation determinant. Adoption is concentrated among documentation sites and developer tooling. It costs little to implement, but it is not a substitute for the more substantial work of indexability, schema, and authority.

The third claim is that AI traffic converts at five to fourteen times the rate of organic search and should therefore receive proportional budget reallocation. The vendor data on this point, principally from Similarweb and from analyses of Microsoft Clarity referrer data, supports a directional finding that AI-referred traffic is higher intent. A peer-reviewed academic study by Schulze and Kaiser, published in October 2025 and analysing 973 e-commerce sites between August 2024 and July 2025, reached a more sober conclusion. It found that ChatGPT-referred traffic accounted for less than 0.2 per cent of total traffic to those sites and underperformed organic search on conversion rate. The honest position is that AI traffic is small but high intent, and that vendors have an incentive to over-claim. We recommend treating it as additive to existing search budgets, not as a reason to redistribute them.

A working definition for SKOMA's clients

For the purposes of this report, we define GEO as follows. Generative Engine Optimisation is the practice of engineering the conditions under which a brand is cited inside the synthesised responses of AI-mediated discovery surfaces. It is a subset of the broader discipline of search, not a replacement for it. Its core levers are entity authority, content density and evidence, semantic structure, technical accessibility, and brand presence on the platforms that LLMs already trust.

EDITORIAL POSITION

GEO is a useful name. It is not a new discipline.

The most productive way to read GEO is as a vocabulary update, not a method update. Almost everything an experienced SEO team should be doing in 2026 was already on the curriculum. What has changed is the relative weighting of the levers, and the interface against which success is measured.

03

How AI Systems Discover Information.

A working understanding of retrieval-augmented generation, embeddings, and knowledge graphs, written for marketers and founders rather than for engineers.

The plumbing, briefly explained.

A working knowledge of how AI systems retrieve information is the difference between a strategy that compounds and a strategy that chases the headline of the month. The mechanics are not as exotic as they are sometimes presented.

An AI search engine, simplified, performs four operations every time a user submits a query. It interprets the query, retrieves a candidate set of relevant chunks of content from an index, selects from that candidate set the chunks that best support an answer, and synthesises those chunks into prose with citations attached. Each of those operations involves engineering decisions that determine which sources rise and which do not.

Step one. Query interpretation.

The user's query is converted into a set of internal representations. In practice this means at minimum that the system creates an embedding, a high-dimensional numeric vector that represents the meaning of the query in a way amenable to similarity search. The embedding model used here is significant. OpenAI uses its text-embedding-3 family. Google uses Gemini Embedding. Anthropic, in its Contextual Retrieval framework, recommends Voyage embeddings paired with BM25. Cohere, BAAI's BGE-M3, NVIDIA's NV-Embed-v2 and others all sit in production at various platforms.

What matters for marketers is not the model name but the consequence. Two queries that share no surface-level keywords can have very similar embeddings if their meaning is similar. A page can be retrieved for a query it does not contain, provided the page's chunks are semantically aligned. This is the technical basis on which keyword-density approaches stopped working some years ago.

Step two. Retrieval.

The system retrieves a candidate set of chunks from an index. There are three common patterns in production. Pure vector retrieval, where chunks are scored by cosine similarity to the query embedding. Lexical retrieval, typically BM25, which still performs surprisingly well on technical and named-entity queries. Hybrid retrieval, which combines both signals, often with reranking by a separate cross-encoder model.

Anthropic published a detailed study of contextual retrieval in September 2024 that is worth understanding. The team showed that the conventional approach of embedding raw chunks loses context. A chunk that reads "this represented an increase of three per cent" does not, on its own, tell the retriever what was increasing or in what period. By prepending a brief, automatically generated context to each chunk before embedding, Anthropic reduced top-twenty retrieval failures by 49 per cent. Adding a reranking step took the reduction to 67 per cent. The headline is that production-grade retrieval is now a multi-step pipeline. Brands that publish chunk-friendly content, with self-contained paragraphs that carry their own context, are favoured by these pipelines.

Step three. Selection.

From a candidate set of perhaps fifty to two hundred chunks, the system must select the small number that will be cited and synthesised into the answer. The selection logic is not fully dis-

closed by any of the major platforms, but published research and platform documentation indicate it draws on at least four signals. Relevance to the query, as measured by the embedding and reranking scores. Authority of the source, often inferred from the source's standing in conventional search rankings or from explicit reputation features. Freshness, with recently updated content scoring meaningfully higher in many domains. Diversity, to ensure the synthesised answer does not repeat the same source three times.

Step four. Synthesis.

The selected chunks are passed to the large language model along with the query and a system prompt that defines how to generate the response. The model produces prose that draws on the chunks, ideally without departing from them, and attaches citations to specific claims. This last step is where citation faithfulness becomes a measurable property. Recent academic work, including a cross-model audit published in early 2026 (Penn et al., arXiv 2603.03299), found that between 3 and 13 per cent of citation URLs produced by commercial LLMs were hallucinated, and a further 5 to 18 per cent failed to resolve to relevant content. Deep research agents, which generate longer responses with more citations, hallucinate at higher rates than search-augmented chat.

The role of knowledge graphs

Sitting alongside the retrieval pipeline at every major platform is a knowledge graph. Google's Knowledge Graph contains hundreds of billions of facts about people, places, organisations, and concepts. It is fed by Wikidata, by Wikipedia, by structured data from the open web (largely Schema.org markup), and by Google's own internal extraction.

The role of the knowledge graph in AI retrieval is to disambiguate. When a user asks about "Tesla", the model needs to determine whether the intended entity is the car company, Nikola Tesla, or the band. Knowledge graphs provide that disambiguation cheaply and reliably. Brands that exist coherently inside the Knowledge Graph and Wikidata are correctly identified and consistently associated with the topics they belong to. Brands that do not exist, or that exist with conflicting or thin entries, are routinely confused, misattributed, or omitted.

This is one of the reasons that establishing a brand's entity presence is among the highest-value moves available to a GEO programme. Wikidata in particular has a lower notability bar than Wikipedia and is, for most B2B brands, the more accessible starting point.

Embedding models and what they imply for content

The Massive Text Embedding Benchmark (MTEB), maintained on Hugging Face, ranks embedding models on retrieval performance across a wide range of tasks. As of April 2026, Voyage-3-large, NV-Embed-v2, BGE-M3, Cohere embed-v4 and OpenAI text-embedding-3-large are the most-deployed in production. The models differ in detail, but they share three properties relevant to content strategy.

First, they reward content that is unambiguous. A paragraph that uses pronouns to refer back to entities mentioned three paragraphs earlier loses information when chunked. Pages that restate their subject explicitly within each section retain meaning when chunked.

Second, they reward content that is internally consistent. A page that presents a single claim with supporting evidence is encoded more cleanly than one that presents three claims interleaved. The implication is not that pages should be short, but that sections within pages should be self-contained.

Third, they reward content that uses the language of the domain. The embedding space for medical content is shaped by the language of medicine. The embedding space for software documentation is shaped by the language of software. Pages that translate their topic into the user's everyday language without reflecting back the domain language can perform less well in retrieval, even if they perform well in user comprehension.

The implication for content systems

The retrieval pipeline above explains, in simple mechanical terms, why the techniques that the Princeton GEO study found effective are effective. Adding citations and statistics is not a stylistic preference of the LLM. It is a signal of evidence that the retrieval pipeline can score, and that the synthesis step can lean on. Adding authoritative quotations is not a flourish. It improves the source-credibility signal that the selection step considers. Front-loading the answer is not a matter of taste. It is a structural advantage in chunk-based retrieval, where the first chunk of a page carries disproportionate weight.

Brands that publish chunk-friendly content, with self-contained paragraphs that carry their own context, are favoured by these pipelines.

CHAPTER THREE

None of this requires a marketer to read research papers. It does require an editorial system that values clarity, evidence, structure, and consistency. Those are not new requirements. They are, in our experience, the requirements that distinguish credible publishing from filler at every level of the discipline. AI surfaces have made the gap between the two more economically consequential than it has been in any previous decade of search.

04

Understanding AI Retrieval and Citations.

How LLMs decide what to cite, why citation faithfulness still trails human standards, and what the practical implications are for sourcing strategy.

What gets quoted, and why.

If retrieval is the engine room, citation is the part the user actually sees. The decisions an AI engine makes about which sources to name, and how, are the decisions that translate technical retrieval into commercial visibility.

Citation behaviour differs across AI engines more than most marketers realise. Google's AI Overviews show, on average, fifteen sources per response in February 2026 measurements, up from under seven a year earlier. Perplexity exposes its citations as the primary affordance of the user interface. ChatGPT cites inline by default and reveals a sources panel at the side. Claude cites with high precision when its citations API is invoked but with much sparser frequency in standard chat. Gemini cites visibly in AI Mode and AI Overviews but more sparingly in conversational sessions. The result is that the same query can produce a different citation footprint on each platform, and the brands that appear differ accordingly.

Where citations come from on each platform

Several independent studies converge on a set of patterns. The Lantern study published in February 2026, drawing on approximately two hundred million citations, found that Reddit accounted for 11.4 per cent of Google AI Overview citations, and that YouTube accounted for 29.5 per cent. Wikipedia consistently appears in the top three across every engine measured. Profound's analysis of 680 million citations between August 2024 and June 2025 found Wikipedia at 7.8 per cent of ChatGPT citations and Reddit at 1.8 per cent. Perplexity weighted YouTube at approximately 16 per cent and Wikipedia at approximately 12 per cent. Each engine has, in effect, learned to trust a slightly different mix of sources.




FIGURE 4.1 Domain citation share across the three largest AI search surfaces.

Source. Lantern, February 2026; Profound, August 2024 to June 2025; Ahrefs analysis of 78.6 million queries. Each platform tracks its own population of queries; comparisons are directional.

Why these sources and not others

Three properties characterise the sources that appear repeatedly across AI engines. They are dense in primary information, structurally consistent, and externally validated.

Wikipedia is dense in primary information because every article is the result of citation-led drafting, with claims tied to external references. It is structurally consistent because Wikipedia's templates, infoboxes, and section conventions make it readable to retrieval systems with very little ambiguity. It is externally validated because thousands of editors review each article over time, and because its contents are mirrored in Wikidata and the Knowledge Graph.

Reddit is dense in primary information for a different reason. It contains first-person experiential testimony at scale, often the only viable source for queries that begin with "is X worth it" or "how did you find Y". It is structurally consistent because every post follows the same template. It is externally validated, in the LLM's eyes, by the Google and OpenAI licensing agreements signed in 2024 that gave both platforms preferential access to the corpus.

YouTube is dense in primary information because video transcripts contain the words spoken in the video, often by domain experts. It is structurally consistent because every video has a title, description, transcript, and channel. It is externally validated through subscriber counts, watch time, and the broader YouTube ranking signals.

The consequence for brands is straightforward. The fastest way for most brands to increase AI citation share is not to publish more on their own site. It is to ensure that the platforms LLMs already trust contain credible representation of the brand and its expertise. A brand that has no Wikipedia entry, no Wikidata record, no LinkedIn presence for its leaders, no Reddit footprint, and no YouTube video assets has chosen, by default, to be invisible in the surfaces that matter most. The choice is reversible, but it requires effort.

The faithfulness question

Not every citation an AI engine produces is accurate. The peer-reviewed and pre-print literature on this question grew rapidly through 2025 and into 2026. A representative finding is that of Penn and colleagues, whose cross-model audit of phantom citations published in early 2026 found that between 3 and 13 per cent of citation URLs produced by commercial LLMs were hallucinated, that a further 5 to 18 per cent failed to resolve to relevant content, and that deep research agents (which produce longer responses with more citations) hallucinated at higher rates than search-augmented chat.

The mitigation that the literature recommends is multi-model consensus. The same study found that when three or more LLMs cite the same source for the same claim, the accuracy of the cited URL rises to 95.6 per cent. The implication is that brands that achieve consistent citation across multiple engines are not just more visible. They are more trusted, both by users and by future AI systems trained on the citations of current ones.

How citations decay or compound

An overlooked feature of AI citation is that, unlike a Google ranking, it is not a real-time market. Citations made in May produce screenshots that circulate in marketing decks in August. Citations from large engines are sometimes re-ingested into the training data of subsequent model versions. The compounding effect is real but slow.

Decay is also slower than in conventional SEO. A page that ranked third for a query a year ago and now ranks tenth produces an immediate traffic decline. A page cited in an AI Overview a year ago, when its content has not been updated, may still be cited today, although freshness

signals nudge against it. Brands that update their cornerstone pages quarterly, with explicit publication and revision dates marked, are favoured by every freshness-aware retrieval pipeline we have evidence of.

Practical implications for sourcing strategy

Three practical positions follow.

Treat external platforms as part of the owned-content estate. A brand's Wikipedia entry, Wikidata record, LinkedIn page, YouTube channel, Reddit presence and G2 listings are not adjacent to the website. For AI visibility, they are at least as material. Treat them with the same editorial care.

Invest in original research at the page level. Pages that publish first-party data, surveys, benchmarks, or experimental results are cited at materially higher rates because they are the only source for their claims. The investment cost is real. The return, in citation share, is durable.

Track citation share, not citation count. A useful KPI is the share of voice your brand achieves across the top one hundred AI prompts in your category, not the absolute number of times it is cited. Tooling for this measurement matured significantly in 2025 and is now within reach of mid-market budgets.

COMMON MISREADING

"Our brand is cited by ChatGPT, so we are doing well."

A single citation is a starting point, not a result. The relevant questions are which queries cite the brand, which competitors are cited alongside it, what the citation says about the brand, and how stable that citation is across model updates. Without those four data points, a citation is anecdotal.

05

AI Search Platforms Compared.

Five engines, five citation logics, five sets of implications. The clearest single page of analysis you will find anywhere on the platform-by-platform shape of the AI search market.

Five engines, five logics.

The temptation to treat AI search as a single category is strong, and wrong. Each major engine has a distinct architecture, a distinct citation logic, and a distinct set of behaviours that reward different kinds of preparation.

The five engines that matter most for general-market visibility in mid-2026 are Google AI Overviews and AI Mode, ChatGPT Search, Perplexity, Microsoft Bing Copilot, and Anthropic's Claude. A second tier of consumer-facing tools includes Brave Leo, You.com, Arc Search and DuckDuckGo Assist. We focus on the first five because they currently account for the overwhelming majority of branded AI search traffic in published measurements.

Google AI Overviews and AI Mode

Architecture. Both surfaces sit on top of Google Search. Both use Gemini, with Gemini 3 becoming default in early 2026. AI Overviews appear at the top of certain SERPs as a synthesised answer with linked source cards. AI Mode is a separate, more conversational interface accessible from a tab in Search. Australia received AI Overviews on 29 October 2024 and AI Mode on 8 October 2025.

Citation behaviour. AIO citations now overlap meaningfully with organic top ten rankings, with BrightEdge measuring 54.5 per cent overlap by September 2025. The average AIO contains fifteen sources, drawn predominantly from YouTube, Reddit, Wikipedia, Quora, and high-authority editorial domains. Healthcare AIOs lean heavily on .gov, NIH, Mayo Clinic and similar authoritative health publishers, with NIH alone holding approximately 60 per cent citation share inside healthcare AIOs.

Implications for brands. The single most consequential investment for AIO visibility remains conventional E-E-A-T-aligned content that ranks organically. Brands ranking outside the top ten are unlikely to be cited in AIO. Brands ranking inside it are increasingly likely to be. AIO is, in this sense, the surface that most rewards continuity of conventional SEO discipline.

ChatGPT Search

Architecture. ChatGPT Search launched in October 2024, was extended to all logged-in users in December 2024 and to all users in February 2025. It is built on Bing's index, augmented by OpenAI's own crawling via OAI-SearchBot and ChatGPT-User. The search interface inside ChatGPT is integrated with conversation memory, so a query may draw on prior conversation context.

Citation behaviour. Wikipedia accounts for the largest single share of ChatGPT citations, around 7.8 per cent. Reddit, Forbes, G2, LinkedIn, and a long tail of authoritative editorial sites follow. The system shows a strong preference for sources that are already trusted by Bing's index. Brands that are absent from Bing or that have weak Bing presence are at a structural disadvantage.

Implications for brands. Most Australian brands have invested less in Bing visibility than in Google. The cheapest visibility increase available in ChatGPT Search is to ensure complete and current Bing Webmaster Tools coverage, IndexNow integration where relevant, and a clean profile in the Bing index. Bing's AI Performance dashboard, launched in February 2026, is now the most useful single source of ground-truth on what ChatGPT-class engines are searching for and citing.

Perplexity

Architecture. Perplexity built its own crawler infrastructure (PerplexityBot for indexing, Perplexity-User for live retrieval) and operates its own ranking and reranking pipeline. The user interface is citation-first by design. Pro Search, Spaces, Internal Knowledge Search, and Deep Research mode form the surfaces of the consumer product.

Citation behaviour. Perplexity weights YouTube and Wikipedia more heavily than the other major engines and surfaces between five and twenty citations per response. The Cloudflare investigation in August 2025 raised concerns about Perplexity using undeclared crawlers and headless browsers; Perplexity disputed those findings. Brands that block PerplexityBot at the WAF or robots.txt layer remove themselves from this surface entirely.

Implications for brands. Perplexity is the engine that rewards depth and specificity most directly. Pages that contain explicit data, named experts, and source citations are noticeably more likely to be cited. The brand that publishes a single piece of original research per quarter, with clear methodology and primary data, will outperform the brand that publishes thirty thinly-sourced articles in the same period.

Microsoft Bing Copilot

Architecture. Bing Copilot is a layer on Bing Search and is now also the underlying retrieval substrate for Microsoft Copilot in 365, Edge, Windows and the consumer Copilot app. The November 2025 redesign introduced more prominent inline citations and a dedicated Search experience inside Copilot.

Citation behaviour. Citations follow Bing rankings closely, with the Bing Webmaster Guidelines updated in February 2026 to reflect Generative Engine Optimisation explicitly. The guidelines now include guidance on noarchive and data-nosnippet for AI controls and have expanded the definition of keyword stuffing to include AI-engineered language patterns.

Implications for brands. The combination of Bing Copilot and ChatGPT Search means that Bing-side visibility is now meaningful for two of the five major AI surfaces. This is a material shift in the economics of investing in Bing optimisation. The Bing Webmaster Tools AI Performance dashboard is the most actionable single source of AI search analytics in the public market.

Anthropic Claude

Architecture. Claude received production web search capability for paid users in March 2025 and an API web search tool in September 2025. Claude operates a three-bot architecture (ClaudeBot for training, Claude-User for user-triggered fetches, Claude-SearchBot for retrieval) which it documents in its developer pages.

Citation behaviour. Claude cites with high precision when its citations API is invoked, with the system designed to minimise unsupported claims. In standard chat, citations are sparser and the model tends to summarise rather than enumerate. Claude is heavily used in enterprise contexts, including by SKOMA's own analysts in the production of this report.

Implications for brands. Claude's user base is smaller than ChatGPT's but skews more enterprise and technical. Brands selling into B2B, technical, financial, healthcare and legal markets will find Claude's audience disproportionately valuable. The recommendation is to treat ClaudeBot as a permitted crawler unless there is a specific reason not to, and to prioritise content that is dense in evidence and explicit reasoning.

The comparison at a glance

ENGINE	INDEX	CITATION STYLE	TOP SOURCES	BRAND IMPLICATION
Google AIO & AI Mode	Google Search	Linked source cards, ~15 per AIO	YouTube, Reddit, Wikipedia, Quora, .gov	Conventional E-E-A-T continues to determine eligibility
ChatGPT Search	Bing & OAI crawl	Inline plus sources panel	Wikipedia, Reddit, Forbes, G2, LinkedIn	Bing-side visibility is the underrated lever
Perplexity	Proprietary crawl	Citation-first UI, 5 to 20 sources	YouTube, Wikipedia, Reddit, primary research	Rewards depth, original data, named experts
Bing Copilot	Bing Search	Inline citations, recently more prominent	Bing-ranked authoritative editorial	Bing Webmaster Tools is the dashboard that exists
Claude	Live web tool, optional API	Sparse in chat, dense via citations API	Authoritative editorial, .edu, .gov, primary docs	Disproportionate value in B2B, technical, regulated

PRACTICAL RECOMMENDATION

Track all five engines. Optimise for the underlying signals.

The temptation is to pick the engine your team has the most enthusiasm for and optimise primarily for that one. The error is that the underlying signals across all five engines are the same. Strong organic ranking, complete entity records, original research, expert authorship, dense and well-structured content, and active presence on Wikipedia, Reddit, YouTube and LinkedIn improve performance everywhere. Single-engine programmes leave compounding returns on the table.

06

The New Visibility Framework.

A five-layer model for thinking about brand presence in AI-mediated discovery, designed to replace the single-axis ranking model that no longer fits the surface.

Five layers of presence.

For twenty years, search visibility was a single number. Position one through ten on a results page. The new visibility surface requires a layered framework, and we offer one here.

The framework that follows is the one SKOMA uses internally. We have refined it through engagements across regulated industries (health, legal, financial), technology, retail, and education. The five layers are presented in order of dependency. Each subsequent layer assumes the layer beneath it is in reasonable condition.

Layer One. Indexability.

The foundation of all visibility, AI or conventional, is whether a brand's content can be retrieved at all. Indexability includes crawlability for the relevant bots, indexed status in Google and Bing, the absence of noindex tags or robots.txt blocks where they are not intended, and the absence of CDN or WAF rules that silently exclude AI search bots.

The most common failures we encounter at this layer are inherited rather than chosen. A WAF rule from 2023 blocks GPTBot. A robots.txt entry written for "AI bots" blocks Claude-SearchBot, which is the bot that retrieves for ChatGPT-style answers in some flows. A site migration leaves canonicals pointing to staging URLs. None of these issues is technically difficult to fix. All of them are common, and any of them can erase a brand from the AI surface for months before being detected.

Layer Two. Entity Identity.

Above indexability sits the question of whether a brand exists, in the eyes of the AI surface, as a coherent entity. The minimum bar is a clean Knowledge Graph entry, a Wikidata record, complete Schema.org Organization markup with sameAs links connecting the website to the brand's social, professional, and reference profiles, and a Wikipedia article where notability supports one.

Brands that fail at this layer are often confused with similarly-named entities, attributed incorrectly, or omitted from answers where they should rightfully appear. Investing here yields disproportionate returns because it improves performance across every AI surface simultaneously.

Layer Three. Authority Signals.

The third layer is the conventional E-E-A-T pillar. Real expertise, demonstrated through named authors with verifiable credentials. Authoritative third-party validation, demonstrated through editorial mentions, citations, and inbound links from reputable sources. A clean reputation, demonstrated through reviews, press, and the absence of policy issues. Trust signals, demonstrated through transparent ownership, contact information, and editorial processes.

Authority is cumulative. It is not produced by a single campaign. The brands that have led at this layer have, almost without exception, invested consistently for at least three years. The brands that begin investing now will see compounding returns in 2027 and 2028. There is no faster path. We have looked.

Layer Four. Content Density.

The fourth layer is what is sometimes called GEO content. Pages that are dense in evidence, structured for chunked retrieval, written in the language of the domain, internally consistent, and updated on a defensible schedule. Original research, primary data, named expert quotations, explicit citations, and clear publication dates are all properties that boost AI citation likelihood.

Content density without the layers below it produces volatile, unreliable performance. Content density on top of indexability, entity identity, and authority produces compounding returns. The order matters.


Layer Five. Distributed Presence.

The top layer is the brand's presence on the platforms LLMs already trust. Wikipedia, Wikidata, Reddit, YouTube, LinkedIn, G2 and the relevant industry publications. This is not a substitute for an owned-content strategy. It is a complement that, in many categories, now produces a larger share of AI citation surface than the brand's own website does.

For brands operating in regulated industries (health, legal, financial, education), distributed presence on industry registries, professional associations, and government data sources is also material. The Australian context is particularly favourable here. Sites like ATAR, Medicare, Health Direct, ASIC, the ATO, and various state government professional registries all carry meaningful weight in AI surfaces because they are already trusted by Google and by extension by the engines that consume Google's index.

The framework, visualised

The diagram below summarises the dependencies. Each layer rests on the one below it. A brand investing only at Layer Four (content) without a healthy Layer One or Two foundation will produce content that does not retrieve, or does not synthesise into an answer because the retrieval pipeline cannot resolve the brand's identity.



SKOMA Visibility Framework. Each layer depends on the one below it.

FIGURE 6.1 The SKOMA Visibility Framework.

SKOMA Digital, 2026. Used in client engagements to diagnose where in the visibility stack a brand has under-invested.

Diagnostic use

The framework's most useful application is diagnostic. When a brand asks why it is not being cited inside ChatGPT or AI Overviews, the productive response is not to write more content. It is to walk the framework from the bottom up. Are we indexed? Do we exist as a coherent entity? Do we have real authority signals? Do our pages express that authority in the form retrieval pipelines reward? Are we present on the platforms that LLMs trust? In our experience, the answer is usually that the brand has invested heavily at one or two layers and almost not at all at one or two others. The deficit, not the strength, determines the result.

The deficit, not the strength, determines the result.

SKOMA ENGAGEMENT PRINCIPLE

07

Entity Authority and Trust.

Why a small set of brands repeats inside almost every AI answer, what those brands have in common, and how to build the same kind of presence over the next two to three years.

Why some brands repeat in every answer.

A small set of brands appears with disproportionate frequency in AI-generated answers. The reasons are not mysterious, and the playbook for joining them, while long, is not complicated.

Across every dataset we have examined, including Lantern's two hundred million citation analysis, Profound's six hundred and eighty million, Ahrefs' seventy-eight million query study, and Semrush's two hundred and thirty thousand prompt research, the pattern is consistent. A small number of brands account for a large share of category citations. In B2B software, G2 and Gartner appear in most categories. In healthcare, Mayo Clinic, Cleveland Clinic, NIH, WebMD and Health Direct (in Australia) recur. In finance, Bloomberg, the Wall Street Journal, Investopedia, the Australian Securities and Investments Commission, and the major bank brands appear repeatedly. In law, FindLaw, the relevant state law society, and a small set of established legal information sites dominate.

What these brands have in common is not size. There are large brands with negligible AI citation share. What they have in common is six properties.

Six properties of repeatedly-cited brands

One. Coherent entity presence in the Knowledge Graph and Wikidata. The brand exists as a single, well-defined entity, with sameAs links to social profiles, business registries, and where applicable Wikipedia. The entity record contains the brand's category, founding date, founders, primary topic areas, and notable products or services.

Two. A library of original primary content. The brand has published, over multiple years, content that is the original source for some specific class of claims. Mayo Clinic's symptom checker. Bloomberg's market data. The Reuters Institute's annual digital news report. Gartner's Magic Quadrants. The brand is not aggregating; it is publishing first-party.

Three. Authoritative third-party validation. The brand is referenced, cited, and quoted by other authoritative sources. This produces a consistent semantic association between the brand and its topics that is detectable to retrieval pipelines and to the synthesis step.

Four. Named, credentialed authors. Content carries the bylines of identifiable individuals with verifiable credentials. The author pages link to LinkedIn, ORCID where applicable, and other authoritative profiles. Schema.org Person markup makes the relationship machine-readable.

Five. Stable, well-maintained URL structure. Pages that have existed for a long time, at the same URL, accruing references, are easier for retrieval pipelines to ground citations against. URL volatility is a quiet killer of AI visibility.

Six. Active presence on the platforms LLMs trust. The brand has a complete LinkedIn footprint, a curated YouTube channel where applicable, a clean Wikidata record, and where the audience is consumer-facing, a credible Reddit and Quora presence.

The Reddit and Wikipedia question

The two platforms that account for an outsized share of AI citation are Reddit and Wikipedia. Both reward presence, but each requires a different kind of presence.

Wikipedia presence requires that the brand cross the notability threshold. For B2C consumer brands, this is usually achieved over time through media coverage. For B2B brands, the bar is higher and the investment more deliberate. The starting point is almost always a clean Wikidata record. Wikidata's notability criteria are looser, and a Wikidata entity is itself a citable, link-worthy reference. We have seen mid-market B2B brands achieve material AI citation lift simply by establishing complete Wikidata records, before any work on the Wikipedia article itself.

Reddit presence is fundamentally different. Reddit communities reject promotional content and reward genuine participation. The brands that succeed on Reddit do so by becoming a credible voice within the relevant subreddits over a period of months and years, with a model that values contribution at roughly a ratio of nineteen to one against any direct promotion. The brands that fail try to shortcut this with marketing-style activity that the community recognises and downvotes within hours.

Reddit's licensing agreements with Google (estimated at sixty million dollars per year, signed in February 2024) and OpenAI (estimated at seventy million dollars per year) have entrenched Reddit content as a primary source for the largest AI surfaces. Brands that ignored Reddit in 2024 and 2025 are now playing catch-up against competitors who built community standing earlier.

How E-E-A-T evolved through the AI transition

Google's Search Quality Rater Guidelines have been updated continuously since the Experience pillar was added in December 2022. The most consequential recent updates are from September 2025, which added explicit AI Overview evaluation criteria and expanded the Your Money or Your Life category to include elections, civic institutions and government trust topics. The September 2025 revision also explicitly elevated Trust as the most important pillar of the four, with the rationale that an unreliable source can undermine an otherwise high-quality result.

The implication is that the older interpretation of E-E-A-T as a content guidance has shifted. It is now better understood as a brand-level diagnostic. The question is not whether a particular page demonstrates expertise. It is whether the brand publishing the page is the kind of brand that Google's quality systems trust to be expert across its category.

Author entity systems

One under-implemented element of authority engineering is author entity. Most websites publish content under generic bylines, an internal team page, or no byline at all. AI surfaces increasingly reward content with named authors who can be resolved as entities and whose credentials are verifiable.

The pattern that works in our engagements is to invest in three to five hero authors per brand. Each hero author has a complete profile page, Schema.org Person markup linking to LinkedIn, ORCID where applicable, and any professional registry, a curated body of work, and external presence on industry platforms (Substack, Medium, conference speaking, podcasts). Content under their byline carries author authority that the brand cannot otherwise produce.

The Australian advantage in regulated industries

Australia's regulated professional industries operate inside a structured information ecosystem that is favourable for AI authority. The Australian Health Practitioner Regulation Agency, the Tax Practitioners Board, the Legal Profession Uniform Law Australian Solicitors' Conduct Rules, the Australian Securities and Investments Commission's professional registers, the Department of Education, and various state-level registries all publish structured data about credentialed professionals. This data is consumed, directly or indirectly, by the same retrieval pipelines that decide what AI engines cite.

An Australian medical practice, law firm, financial advice business, or accredited training organisation that ensures complete and accurate listings on its relevant registry, with consistent NAP (name, address, phone) data and complete sameAs entity links, is establishing a layer of trust that international competitors cannot easily replicate. This is one of the genuinely under-priced opportunities in the Australian market today.

AUSTRALIAN PRACTICE NOTE

Audit your registry footprint before your content footprint.

For regulated Australian businesses, the first ninety days of any AI authority programme should include a complete audit of how the business appears across its industry registries, professional associations, and government data sources. Inconsistent or incomplete registry data is a direct signal of weak entity coherence to AI surfaces. The fix is administrative, not editorial, and it produces visibility lift across every AI engine simultaneously.

08

Content Systems for AI Discovery.

What the Princeton GEO study actually proved, what the follow-on research has refined, and how to translate the academic record into editorial standards your team can apply tomorrow.

What the studies actually show.

More commentary has been written about content optimisation for AI surfaces than about any other aspect of GEO. Most of it is downstream of a single academic paper. It is worth reading the paper carefully before drawing conclusions.

The Princeton, Georgia Tech and IIT Delhi paper, formally titled *GEO: Generative Engine Optimization*, was first posted to arXiv in November 2023 and presented at the ACM SIGKDD KDD conference in Barcelona in August 2024. It is the most-cited primary source in the GEO field and remains, twenty months later, the most rigorous published analysis we have access to.

The study tested nine content modifications across ten thousand queries, using five generative engines. The modifications fell into three categories. Stylistic, including authoritative tone and fluency optimisation. Structural, including bullet points and statistics addition. And source-related, including cite sources, quotation addition, and adding URLs to content.

What the study found

The headline finding is that source-related modifications had the largest and most consistent positive effect on AI citation visibility, with average lifts in the range of thirty to forty per cent. Adding statistics, citing sources, and adding direct quotations from authoritative figures each produced a comparable lift in the average across queries. Fluency optimisation produced a smaller but still positive effect of around fifteen per cent. Authoritative tone produced an eight per cent average lift but was effective only in debate-style and historical queries, with null or negative effects in factual queries. Keyword stuffing produced a negative effect of approximately ten per cent, the only modification in the set that worsened performance.




FIGURE 8.1 Effect of nine content modifications on AI citation visibility, averaged across ten thousand queries.

Source. Aggarwal et al., *GEO: Generative Engine Optimization*, ACM SIGKDD KDD 2024. Coral bars indicate the source-related modifications that produced the largest lifts.

Domain effects

Aggregating across all ten thousand queries, the source-related modifications dominate. But the domain-level breakdown is also informative. In factual queries, citing sources produced the largest lift. In debate queries, citing sources produced the largest lift. In debate queries, quotations from authoritative figures led. In historical queries, both quotations and citations performed strongly.




FIGURE 8.2 Effect of three source-related modifications across query domains.

Source. Aggarwal et al., 2024. The lift is consistent across domains but the relative ordering shifts.

What follow-on research has refined

Several follow-on studies in 2025 and early 2026 refined and qualified the original finding.

First, position within the page matters. ALM Corp's analysis of ChatGPT citations found that 44 per cent of citations were drawn from the first third of the cited page. The implication is that the answer should not be in the conclusion. It should be near the top of the page, structured for chunk-based retrieval.




FIGURE 8.3 Where on a cited page LLMs draw their citations from.

Source. ALM Corp citation position study, 2025. Sample of cited pages across ChatGPT responses.

Second, freshness compounds with evidence. ConvertMate's GEO benchmark study, published in early 2026, found that pages updated within the previous thirty days were cited at approximately three times the rate of equivalent pages older than six months. The lift was independent of the original Aggarwal effect, suggesting that freshness and evidence are additive levers.

Third, the model-by-model variation is meaningful. The Aggarwal study tested five engines and found that the average lift held across them, but the magnitude varied. The follow-on work suggests that source-related modifications produce the largest lift in Perplexity, the second-largest in Google AI Overviews, and a smaller but still meaningful lift in ChatGPT.

From study findings to editorial standards

Translating the academic record into a set of editorial standards your team can apply produces something like the following. We have used these standards inside SKOMA engagements with measurable effect.

EDITORIAL STANDARD	SPECIFICATION	EVIDENCE
Front-loaded answer	The article's main answer in the first 200 to 300 words	44 per cent of LLM citations from first third of page (ALM Corp, 2025)
One credible primary citation per claim	Every non-trivial claim links to a verifiable source	+30 to +40 per cent visibility lift (Aggarwal et al., 2024)
Statistic per major section	Each major section contains one explicit, sourced number	+30 to +40 per cent visibility lift (Aggarwal et al., 2024)
Direct quotation from named authority	One named-expert quotation per article minimum	+30 to +40 per cent visibility lift (Aggarwal et al., 2024)
Defensible publication and revision dates	Visible date stamps; review at least quarterly	~3x citation likelihood for pages updated within 30 days (ConvertMate, 2026)
Explicit author byline	Named author with linked credentials and Person schema	Author entity is consumed by retrieval pipelines
Self-contained section headers	Every section heading restates the topic; pronouns resolved	Chunk-friendly retrieval; explicit in Anthropic Contextual Retrieval, 2024
Schema.org coverage	Article, FAQPage, Person, Organization markup as relevant	Grounding signal across all major AI surfaces

What does not work

The same body of research has equal value as a list of techniques that do not work. We have observed four common content-system failures in the market.

The first is keyword stuffing in modern dress. Aggarwal's negative effect for keyword stuffing is now well understood, but variants persist. Pages that repeat brand or category terms unnaturally to "match" search intent tend to underperform. Retrieval pipelines reward semantic richness, not lexical density.

The second is generic AI-written content at scale. Google's March 2024 core update redefined scaled content abuse to apply regardless of whether content was produced by humans, AI, or a combination. The brands that built large content libraries with generic AI-generated articles in 2023 lost between sixty and ninety-five per cent of their traffic in the affected period, with HouseFresh and similar independent reviewers among the most documented casualties.

The third is over-formatted content that contains no original information. FAQ schema, table-of-contents widgets, hero modules, and accordion-collapsed sections add navigational chrome but do not add evidence. Retrieval pipelines reward the evidence; users reward the chrome only when there is something behind it.

The fourth is undifferentiated topical clusters. The classic SEO playbook of building wide topical coverage across a category remains useful, but it is no longer sufficient. The brands that succeed at the AI surface combine breadth with depth: a few cornerstone pages that contain genuinely original analysis, surrounded by a coherent network of supporting pages. Coverage

without depth produces a thin, low-citation pattern that the AI surface increasingly de-prioritises.

Front-load the answer. Cite the source. Quote the expert. Show the date. The whole framework is essentially that compact.

EDITORIAL DISTILLATION

09

Technical SEO in AI Search.

The crawlers, the schema, the bot policy you have not audited, and the technical foundations that determine whether any of the editorial work above can take effect.

The technical foundations.

Technical SEO has not become less important in the AI era. It has become more consequential, because a misconfigured robots.txt or WAF rule can now erase a brand from the AI surface in a way that conventional ranking would mask.

This chapter walks through five technical topics that, in our experience, deserve direct attention from any team operating an AI visibility programme. The list is not exhaustive but covers the elements where we have seen the largest gap between common practice and what the data supports.

Crawler ecosystem

The number and variety of AI-related crawlers operating in 2025 and 2026 is materially greater than at any point in the prior decade. Cloudflare's annual crawler audit identifies, among the largest AI bots, GPTBot at 30 per cent of AI crawler share, Meta-ExternalAgent at 19 per cent, ClaudeBot at 5.4 per cent, Bytespider at 7 per cent (down from 42 per cent in 2024), and PerplexityBot, Google-Extended and others. ChatGPT-User specifically grew at over 2,800 per cent year on year, indicating that user-triggered live retrieval (as distinct from training crawl) has become the dominant pattern.




FIGURE 9.1 Share of AI crawler traffic by bot, 2025.

Source. Cloudflare 2025 crawler audit. Data covers verified AI bot traffic across the Cloudflare network.

The three-bot architecture

The dominant pattern among the major AI providers is a three-bot architecture, where each provider runs separate bots for distinct purposes. This is critical to understand because it changes the consequences of bot-blocking decisions.

OpenAI runs GPTBot for training data collection, OAI-SearchBot for indexing the open web for ChatGPT Search, and ChatGPT-User for user-triggered fetches. Blocking GPTBot prevents training use of your content. Blocking OAI-SearchBot removes you from ChatGPT Search citations. Blocking ChatGPT-User prevents users from pasting your URL into ChatGPT and having the system retrieve it.

Anthropic runs ClaudeBot for training, Claude-SearchBot for retrieval-augmented generation answers, and Claude-User for user-triggered fetches.

Perplexity runs PerplexityBot for indexing and Perplexity-User for live retrieval.

Google runs Googlebot (which is the conventional crawler also used for AI Overview eligibility), Google-Extended (which controls training data use for Gemini), Google-CloudVertexBot for Vertex AI grounding, and others.

The decision to block a training bot is a defensible position on copyright and brand protection. The decision to block a search or user bot is, in almost every case we have evaluated, an unintended consequence of using a wildcard rule that catches all "AI" bots. The fix is to differentiate explicitly. A robots.txt that allows OAI-SearchBot, Claude-SearchBot, Perplexity-User and Googlebot, while disallowing GPTBot, ClaudeBot, PerplexityBot and Google-Extended, is a coherent position. A robots.txt that disallows all AI bots is a position, but it is more often the result of inattention than of strategy.

The hidden bot-blocker problem

Beyond robots.txt, content delivery networks and web application firewalls increasingly include managed rule sets that block AI crawlers automatically. Cloudflare introduced its "Block AI Bots" toggle in 2024. Akamai, Imperva and others have similar features. These rules sometimes activate by default, and they are not always reflected in the published robots.txt.

From the audits we have conducted, a meaningful share of B2B SaaS sites have at least one AI bot blocked at the WAF or CDN layer that the site owner did not intentionally block. The proportion is higher among sites that adopted aggressive bot-management defaults during 2024. The audit is mechanical: review the WAF rule set, check the published robots.txt, verify with logs that the relevant bots are actually being served 200 responses. Most of the time, no rebuild is required, just a small set of explicit allow rules.

Schema.org as the grounding language

The role of Schema.org has shifted in the AI era. It used to be primarily about rich-result eligibility on the SERP. It is now primarily about providing AI surfaces with unambiguous structured signal about entities, relationships, dates, prices, and credentials.

The schema types that produce the largest practical effect on AI grounding include Organization (with sameAs links to Wikidata, LinkedIn, Crunchbase, Bloomberg and the relevant industry registries), Person (for credentialed authors and executives), Article (with author, datePublished, dateModified, and citation properties), Product (with offers, aggregateRating, and brand), FAQPage (which is read by AI surfaces even though Google has reduced rich-result coverage of it), HowTo (for the same reason), and LocalBusiness (which connects to Google Business Profile, Apple Business Connect, and other registry surfaces).

Schema is not a direct ranking factor by Google's stated position. It is, however, a direct grounding signal for AI retrieval pipelines, and the difference between brands that have implemented schema fully and those that have not is detectable in citation share. Schema App's 2025 analysis

is quantified the lift, with brands that implemented Article and Person schema with author entity linking seeing measurable AI citation increases within two quarters.

Multimodal indexing

AI surfaces consume more than text. Video, images, PDFs and audio are all now indexed and surfaced. YouTube's status as the single most-cited domain in Google AI Overviews is partly a function of the depth of YouTube's transcript and metadata indexing. Brands that publish video content without transcripts, captions, structured chapters and clean metadata are publishing assets the AI surfaces can see but cannot read deeply.

PDFs deserve specific mention because they are the format most commonly under-treated. Bing's AI grounding actively reads PDF content. Google does as well. PDFs that are scanned images rather than text-extractable, that contain no metadata, or that lack the visible publication and revision dates that the editorial standards above require, are visible in search but invisible in citation.

llms.txt and IndexNow

Two protocols deserve a final mention because they appear repeatedly in vendor commentary.

llms.txt is a Markdown-based proposal for publishing a curated index of a site's most important pages, drafted by Jeremy Howard at Answer.AI in 2024. It is implemented by some documentation sites and tooling. It is not used by any of the major LLM platforms as a primary retrieval signal. The cost of implementing it is low. The expected return is also low. Treat it as a hygiene item, not a strategic move.

IndexNow is a Microsoft-led freshness signalling protocol, supported by Bing and Yandex, that allows publishers to push update notifications to indexers. For sites with high publishing cadence, IndexNow has measurable benefit on Bing freshness, which in turn affects ChatGPT Search and Bing Copilot. The implementation cost is low and the return is more clearly evidenced.

AUDIT PRIORITY

Before any new content investment, complete a five-point technical audit.

One. robots.txt review with explicit per-bot decisions. Two. WAF and CDN rule audit for hidden AI bot blocks. Three. Schema.org coverage map across Organization, Person, Article, Product, FAQPage. Four. Multimodal indexing review (video transcripts, image alts, PDF text-extractability). Five. Indexed-status verification across Google, Bing, and the major AI search panels in Profound, Otterly, or Bing Webmaster AI Performance. The audit is one to two weeks of work. The visibility lift it produces is, in most cases we have evaluated, larger than the next ninety days of new content investment.

10

Human Expertise versus AI Scale.

Why the quality bar in publishing has risen, not fallen, with the arrival of generative AI, and how the brands that win the next cycle will use AI as augmentation rather than replacement.

The quality bar, raised.

The arrival of generative AI in 2022 and 2023 produced an early bet that the marginal cost of content would collapse and that whoever published the most would win. The bet did not work. The brands that built the largest libraries of generic AI content lost the most ground.

The data on this question has accumulated rapidly. Google's March 2024 core update, expanded in November 2024 and again in January 2025, redefined scaled content abuse to apply regardless of whether content was produced by humans, by AI, or by a combination. The September 2025 update to the Search Quality Rater Guidelines went further, defining content as "Lowest" only when it lacks human oversight and review. The framework that has emerged is not anti-AI. It is anti-unreviewed.

The effect on the publishing economy has been structural. Independent niche sites that had relied on programmatic AI content lost between sixty and ninety-five per cent of their traffic in the affected period. Recovery has been slow and partial, and limited largely to sites that can demonstrate first-hand expertise.

The hallucination problem and why it matters for publishers

The reliability gap between AI-generated and human-reviewed content remains real. The peer-reviewed literature places hallucination rates for general LLM responses in the range of fifteen to twenty-seven per cent, depending on the model, the task, and the year of the study. The cross-model audit by Penn and colleagues, published in early 2026, found that between three and thirteen per cent of citation URLs produced by commercial LLMs were hallucinated, and that deep research agents were systematically worse than search-augmented chat.

For publishers, the implication is that any AI-generated content that has not been reviewed by a domain expert carries a risk of factual error in the range of three to thirteen per cent at the citation level, and a higher rate at the substantive claim level. The risk is not theoretical. It is the kind of risk that triggers regulatory action in financial advice, that exposes legal practices to professional negligence claims, and that produces the public corrections that compound into reputational damage over time.

The human-in-the-loop model that is winning

The publishers, agencies, and in-house teams that have adapted most successfully share a common pattern. AI is used as a research, drafting, structuring, and editing assistant. A human expert with relevant credentials reviews, revises, and signs off. The published content carries a named byline. The publication and revision dates are explicit. Where AI has materially contributed to content, that contribution is disclosed.

The Reuters Institute's 2025 Generative AI and News Report, drawing on a survey of approximately twelve thousand respondents across six markets, found that only four to nine per cent of weekly news consumers in the surveyed countries were using AI chatbots as a primary news source. Trust in human-produced news remained significantly higher than trust in AI-produced news, with a gap of approximately twenty percentage points. The audience demand for human voice and editorial accountability is real and persistent.

The Edelman 2025 Trust Barometer Flash Poll, surveying five thousand respondents across five markets in October 2025, found a similar pattern. Trust in business to use AI responsibly was higher than trust in government to do the same (forty-six per cent versus thirty-four per cent), and trust in one's own employer to use AI responsibly was the highest single category at fifty-six per cent. The implication is that brands that are seen to use AI responsibly accrue trust, and brands that are seen to use AI carelessly accrue distrust.

How respected publishers are integrating AI

The major publishers that have signed AI licensing agreements with OpenAI, Google, Microsoft and Anthropic include Reuters, Axel Springer, the Associated Press, Le Monde, Condé Nast, Time, Dotdash Meredith, the Financial Times, Vogue, and others. Each of these agreements includes editorial governance provisions. Each requires human oversight clauses. None of them treats AI as a replacement for journalists.

The Independent, the United Kingdom-based publisher, launched its AI-assisted "Bulletin" product in 2024 with a five-to-ten-bullet summary format, with explicit editor sign-off on every bullet. The product is positioned as augmenting rather than replacing the underlying journalism. Substack reported approximately forty per cent year-on-year growth in the twelve months to July 2025, suggesting that audience demand for human-voice, individually-attributed content is, if anything, increasing rather than diminishing.

Practical principles for content systems

Three practical positions follow from the data above. We have used them in client engagements with measurable effect.

Pair every published piece with a credentialed reviewer. The reviewer's role is not editorial nicety. It is professional accountability. In regulated industries, the reviewer is often legally required. In unregulated industries, the reviewer is what distinguishes content that the AI surface considers expert from content that the AI surface considers generic.

Disclose AI involvement transparently. A short editorial standard explaining how AI is used, who reviews it, and what the editorial process is can sit at /editorial-standards or in the footer. Disclosure is now beginning to function as a trust signal in the same way that author bylines and publication dates do.

Reduce the volume. Publishing fewer, better, more authoritatively-attributed pieces is the consistent pattern among brands that have grown AI citation share. Publishing more is the consistent pattern among brands that have lost it. The relationship between volume and authority has reversed, and the publishing systems that have not adapted to this reversal are losing ground that will not return without a structural shift.

The AI content gold rush is over. The expert content era has begun.

FINDINGS, FIVE

11

What Businesses Are Getting Wrong.

Eleven mistakes that the data already records, drawn from agency engagements, public case studies, and the published failures of the early AI content cycle.

Eleven mistakes the data records.

A pattern of recurring errors has emerged across the AI search transition. They are not exotic. They are the predictable consequences of moving fast on incomplete information. The list below is intended as a diagnostic, not as criticism.

One. Blocking the wrong bots.

The most consequential single mistake we encounter is brands blocking the search and user-fetch bots that drive AI citations, while attempting to block the training bots that drive AI training. The result is that the brand exits the AI surface entirely without securing the copyright protection it intended.

Diagnosis: review robots.txt and WAF rules with explicit per-bot decisions. The training bots (GPT-Bot, ClaudeBot, Google-Extended) and the search and user-fetch bots (OAI-SearchBot, Claude-SearchBot, Claude-User, Perplexity-User, Googlebot) are different. Treat them differently.

Two. Treating llms.txt as a silver bullet.

The llms.txt proposal has been over-claimed in vendor commentary. No major LLM platform uses it as a primary retrieval signal. Implementing it is a low-cost hygiene item. Treating it as a substitute for indexability, schema, or authority is a misallocation of attention.

Three. Optimising for one engine and assuming it transfers.

Each of the five major AI engines uses a different citation logic. Profound's analysis of 680 million citations found only eleven per cent overlap in top sources between ChatGPT and Perplexity. A brand that achieves ChatGPT visibility and assumes Perplexity will follow is making an unsupported inference. Multi-engine measurement is now the baseline.

Four. Measuring success only by clicks.

If clicks remain the only KPI, AI-mediated visibility produces apparent decline even when brand presence is rising. The metrics that now matter alongside clicks are citation share, brand mention rate, sentiment within AI answers, and share of voice across the prompts that map to commercial intent in the brand's category. Tooling for these metrics matured in 2025 and is now within reach of mid-market budgets.

Five. Treating GEO as a formatting exercise.

Adding FAQ schema, bullet lists, and table-of-contents widgets without changing the underlying authority and evidence of the content produces volatile, unreliable results. The Princeton GEO study showed that the lift comes from citing sources, adding statistics, and adding authoritative quotations. Those are content investments, not formatting changes.

Six. Producing AI content at scale.

The case studies are now public. HouseFresh, Retro Dodo, Travel Lemming, Charleston Crafted, GGRecon, Test Coches, and others lost between sixty and ninety-five per cent of their traffic in the affected period. Even brands that produced AI content with light editorial review saw sustained ranking suppression. Google's quality systems treat content produced at scale as a quality signal, regardless of method. Volume is no longer a defensible strategy.

Seven. Ignoring Reddit, Wikipedia, and YouTube.

The platforms LLMs already trust are the platforms most brands have not yet treated as part of their owned-content estate. Reddit alone accounts for over eleven per cent of Google AI Overview citations in some studies and over six per cent of Perplexity citations. Wikipedia is consistently in the top three across every engine. YouTube takes nearly thirty per cent of Google AIO citations. Brands that under-invest in these surfaces are under-investing in AI visibility, full stop.

Eight. No entity work.

Brands without coherent Knowledge Graph entries, complete Wikidata records, full Schema.org Organization markup with sameAs links, or named-author Person schema are routinely confused, misattributed, or omitted in AI surfaces. The fix is administrative rather than editorial, and it produces visibility lift across every AI engine simultaneously.

Nine. Weak schema implementation.

Schema is now the grounding language of AI retrieval. Brands that have implemented partial schema (Article without author, Product without offers, FAQPage without question entities) have left the most consequential single technical lever under-deployed. A complete schema audit, with each major content type mapped to its canonical type, is a one-week project that produces sustained citation improvement.

Ten. No author authority.

Anonymous content, generic team bylines, or content attributed to "the editorial team" produces weaker citation signal than content with named, credentialed authors. AI surfaces increasingly reward content that ties claims to identifiable expertise. Three to five hero authors per brand, each with a complete profile and external presence, is the pattern that works.

Eleven. Panic budget reallocation.

The most expensive mistake we have observed is brands shifting fifty per cent or more of their search budget to "AI optimisation" on the basis of vendor projections, then losing organic traffic and not gaining proportional AI visibility. AI traffic is high intent but small in absolute share for most brands. The honest position is that AI visibility is additive to existing search investment, not a replacement for it. The brands that have grown AI citation share have done so by extending their existing programmes, not by abandoning them.

A NOTE ON PATIENCE

The compounding period for authority is two to three years.

The brands that are visible inside AI surfaces today began the relevant work in 2022, 2023, and 2024. The brands that begin in 2026 will see compounding returns in 2027 and 2028. The strategies that produce results in two months are also the strategies that produce ranking suppression six months later. Patience, in this discipline, is now an underrated competitive advantage.

12

AI Workflow Systems for Modern Marketers.

How to integrate AI into editorial, research, and analytical workflows without compromising the quality and accountability that AI surfaces increasingly reward.

Editorial governance for the AI-assisted team.

The teams that produce credible AI-augmented work share a small set of principles. They are not technically demanding. They require clarity about the role of AI, about who is accountable for the output, and about how the work will be reviewed.

This chapter outlines the workflow patterns that we have seen produce reliable, defensible content at scale. The intent is not to recommend specific tools, although we mention several that we use ourselves. The intent is to describe the operating model that translates AI capability into trustworthy publishing output.

The four-stage editorial workflow

The pattern that recurs across high-performing teams has four stages.

Stage one. Research and structure. AI is used to gather evidence, to synthesise prior work, to draft outlines, and to identify gaps. The output is an annotated brief, not a draft. The brief identifies the key claims to be made, the sources that support them, and the structural form the piece should take.

The tools that have matured in this stage include Perplexity Spaces for shared research environments, NotebookLM for source-bounded synthesis where the team uploads its own primary documents, ChatGPT with Projects and Memory for ongoing research threads, and Claude with Projects for long-form analytical work. Each of these tools is most valuable when scoped to specific source material rather than allowed to draw on open-web retrieval alone.

Stage two. Drafting. AI assists in producing a first draft from the annotated brief. The draft is treated as raw material, not as the publication output. The draft is generated, reviewed, and substantially revised by a human author with the expertise to extend, correct, and contextualise.

Stage three. Expert review. A subject-matter expert with relevant credentials reviews the revised draft. In regulated industries, the reviewer is often legally required. In unregulated industries, the reviewer is what distinguishes content that the AI surface considers expert from content that it considers generic. The reviewer's role is not editorial polish. It is professional accountability.

Stage four. Publication and disclosure. The piece is published with a named byline, an explicit publication date, and a brief AI-disclosure statement where AI has materially contributed. The disclosure is a trust signal. The byline is a brand asset. The date is a freshness signal that retrieval pipelines actively reward.

The governance documents that matter

Three lightweight documents anchor a credible AI-assisted editorial system.

An editorial standards page. Public, durable, hosted at a stable URL like /editorial-standards. Describes the editorial process, the role of AI in that process, the credentials and review responsibilities of the editorial team, and the disclosure conventions used when AI has materially

contributed to a piece. This page is read by AI surfaces and by users; both increasingly reward its presence.

A source policy. An internal document specifying which sources the team treats as primary, which as secondary, and which as inadmissible. The policy clarifies the standards for citation, for paraphrase versus direct quotation, and for verification. This document is not public-facing but determines the quality of every piece the team produces.

An author entity programme. A small set of three to five hero authors per brand, each with a complete profile page, complete Person schema, external presence on LinkedIn, ORCID where applicable, and any professional registry, and a curated body of work both on the brand's site and externally. The programme is funded as a business asset, not as a personal-brand favour.

Workspace tools and the pattern of use

The AI workspaces that have proven most useful in editorial settings share a common feature. They allow the team to scope retrieval to specific source material. Perplexity Spaces, Notebook-LM, ChatGPT Projects, Claude Projects, and Gemini Notebooks all support this pattern. The team uploads or links the primary documents that bound the research, and the AI's responses are constrained to that source material.

This is a structurally different use of AI from open-web chat. Open-web chat is useful for ideation, for surfacing perspectives, and for breaking through writer's block. Source-bounded chat is useful for production work where accuracy matters and where the cost of hallucinated claims is high. Mature editorial teams use both, but distinguish them carefully.

Memory systems and what they imply

The memory features rolling out across ChatGPT, Claude, and Gemini through 2024, 2025, and 2026 introduce a new dimension. ChatGPT Memory, expanded in April 2025 to reference all prior conversations. Claude Memory, launched in September 2025 for Team and Enterprise, expanded to Pro and Max in October 2025, and to free users in March 2026. Gemini Personal Context, automatic from August 2025 in most markets and excluded from the European Union, United Kingdom and Switzerland for GDPR reasons.

For editorial workflows, memory features mean that ongoing engagements with an AI assistant accumulate context over time. The team's brand voice, source preferences, editorial standards, and prior research can be referenced without re-explanation. Used well, this compounds productivity. Used badly, it produces drift, where the assistant carries unintended context into new work.

The discipline is to scope memory deliberately. Each major project receives its own workspace or project. Cross-project memory is reserved for genuinely durable elements like brand voice and style guide. The administrative cost of this discipline is small. The productivity benefit, in our experience, is substantial.

The agency model, evolving

For agencies and external partners, the workflow implications are similar but with one additional dimension. The agency is increasingly the operator of editorial governance for the client, not just the producer of content. The client expects, or should expect, that the agency can articulate its AI-use policy, that the agency's named authors carry credentials commensurate with the topic, that the agency's editorial standards are documented and durable, and that the agency can demonstrate the chain of review that backs every published piece.

SKOMA's own editorial standards page is published at skomadigital.com.au/editorial-standards the same reasons we recommend the practice to clients. The standard the discipline now requires is one of demonstrable accountability. The brands that adopt this standard early will find that the same investment serves both their AI visibility and their audience trust.

The standard the discipline now requires is one of demonstrable accountability.

EDITORIAL OPERATIONS

13

Future of Search and Discovery.

Confirmed trends, emerging trends, and speculative ones, separated explicitly. The next two to five years, with calibrated confidence.

The next two to five years.

A great deal has been written about the future of search. We separate confirmed trends from emerging ones, and emerging ones from speculative ones. The discipline of saying which is which is, in our view, the discipline that produces actionable foresight.

Confirmed

Agentic browsing is real, public, and benchmarked. OpenAI's Operator scored 87 per cent on the WebVoyager benchmark when launched in January 2025 and was integrated into ChatGPT agent in July 2025. Google's Project Mariner scored 83.5 per cent and is now embedded in the Gemini Enterprise Agent Platform announced at Cloud Next 2026. Anthropic's Computer Use, which underpins Claude's agentic capabilities, scored 56 per cent. Reliability still trails human performance materially (humans score around 72 per cent on OSWorld, computer-use agents around 38 per cent), but the direction is clear.

The Model Context Protocol has become an industry standard. MCP, originated by Anthropic in late 2024, was donated to the Linux Foundation in February 2026. As of early 2026 there are over ten thousand public MCP servers and approximately ninety-seven million monthly SDK downloads. The Agent-to-Agent (A2A) protocol from Google, announced at Cloud Next 2026, is now in production at over one hundred and fifty organisations. Together these protocols define how AI agents will increasingly interact with brand assets, APIs, and data.

Memory systems are shipping at scale. ChatGPT Memory, Claude Memory across all paid and free tiers, and Gemini Personal Context (auto-on in most markets) all entered general availability in 2025 and 2026. The implication for brands is that AI assistants will increasingly carry context about their users across sessions, and the brand presence within that personalised retrieval becomes a new competitive surface.

Multimodal search is mainstream. Google Lens reached approximately twenty-five billion queries per month by mid-2025. Pinterest Lens, Snapchat's visual search, and the camera-first behaviour of TikTok all point in the same direction. Search Live, the live-video and voice variant of Project Astra, is in active rollout. Visual search via cameras is now the entry point for a non-trivial share of commercial intent, particularly in retail, travel, and home goods.

Emerging

Conversational commerce. ChatGPT shopping, Microsoft Copilot's transactional capability, Google's product embeddings inside AI Overviews, and Perplexity's shopping integrations have all moved from proof-of-concept to live in 2025 and 2026. The volume is small relative to conventional commerce but growing quickly. The brands that have prepared their product feeds, structured pricing data, and merchant integrations for AI ingestion are in an advantaged position when the volume reaches commercial scale.

AI Mode-style interfaces are spreading. The pattern of multi-step, follow-up-aware, conversational search is now visible in Google AI Mode, Bing Copilot's redesigned search, Perplexity's Pro Search, and ChatGPT Search. Engagement metrics on these interfaces are materially higher than on conventional search (queries are two to three times longer, follow-up rates approach

twenty-five per cent). The brands that have written content for the multi-step query, where the user's second and third questions extend the first, will be increasingly visible.

Browser-level AI agents are competing. The Browser Company's sale to Atlassian for approximately six hundred and ten million dollars in September 2025, the launch of Dia (the successor browser absorbing Arc's features), Perplexity Comet, ChatGPT Atlas, Brave AI, Edge Copilot, and Claude for Chrome have established the AI-first browser as a real category. Adoption is concentrated in early-adopter audiences, but the trajectory points towards a meaningful share of high-intent search activity migrating from search engines to in-browser AI assistants over the next two to three years.

Outcome-based content licensing. The Reddit licensing agreements with Google and OpenAI, the rise of TollBit and ProRata, the formation of the Really Simple Licensing coalition, and Cloudflare's bot pay-per-request proposal all point to an emerging market for content licensing that pays publishers for their contribution to AI training and retrieval. The model is early. The direction is clear.

Speculative

Disintermediation of branded websites by agentic browsers. The hypothesis that AI agents will replace direct visits to brand websites for a meaningful share of high-intent transactions is plausible. The current evidence is that agent reliability is still well below human performance, and that for high-stakes decisions users continue to verify directly. The trajectory may produce material disintermediation in three to five years. It is not credible to plan around it as a near-term certainty.

Decentralised discovery. Brave's on-chain top-level domain experiments, the Really Simple Licensing initiative, and various blockchain-based content provenance schemes constitute a real but niche thread. They may become consequential. They have not yet.

Autonomous agentic SEO. The hypothesis that AI agents will themselves transact with sites on behalf of users, requiring brands to optimise for agent comprehension as well as human comprehension, is technically possible. The implementation requires both reliability improvements that have not yet occurred and a level of user trust that has not yet been established.

What we believe is most likely to matter

If we had to nominate the three trends most likely to materially affect brand visibility in the period 2027 to 2030, the list would be the following.

First, agentic browsing reaches consumer scale for navigational and lower-stakes transactional queries. Brands need to be readable to agents, with clean structured data, predictable URLs, and machine-friendly conversion flows.

Second, on-device personal assistants produce a new visibility surface that is not measured by any of the current AI search analytics tools. Brands that achieve presence within this surface will do so through device-level integrations, voice presence, and the same authority signals that drive conventional AI citation.

Third, content licensing becomes a meaningful revenue stream for publishers and a meaningful cost for AI platforms. The economics of the open web shift, with smaller publishers compensated for AI use of their content. The brands that have built credible publishing operations will benefit. The brands that have relied on aggregating others' content will face new structural pressure.

FORECASTING DISCIPLINE**Confirmed, emerging, speculative. Distinguished explicitly.**

The discipline of distinguishing confirmed from emerging from speculative trends is, in our experience, the most useful single habit a strategy team can adopt during the AI search transition. The temptation is to treat all directional trends as equally actionable. The reality is that confirmed trends should be in this quarter's plan, emerging trends should be in this year's plan, and speculative trends should be in the watch-list. Mixing them produces strategic chaos.

14

Practical Action Framework.

A six-step plan for the next ninety days, sequenced by dependency, designed to be executable without specialised tooling and to produce measurable visibility improvement within two quarters.

A six-step plan for the next ninety days.

The framework below is the one SKOMA uses with new clients in the first ninety days of an AI visibility engagement. It is sequenced by dependency. Each step assumes the steps before it are complete.

The total elapsed time for a competent team is twelve to fourteen weeks. The work is substantial but not technically demanding, and most of it is administrative or editorial in nature rather than developmental. The visibility lift, in our engagement data, becomes measurable within two quarters and continues to compound through the second and third quarters.

01

Audit and Measurement Baseline

Weeks 1 to 3. Establish the dashboard before the work. Connect Bing Webmaster Tools AI Performance, Google Search Console AI Overview impression data, and at least one independent AI visibility tracker. Map the top one hundred prompts in your category. Record current citation share by prompt and by engine. Without this baseline, no subsequent improvement can be attributed.

02

Technical Foundations and Bot Policy

Weeks 2 to 5. Audit robots.txt with explicit per-bot decisions across the eight to twelve AI bots that matter. Audit the WAF and CDN rule set for hidden blocks. Verify indexed status across Google, Bing, and the major AI panels. Address any deprecated bot strings. The output is a documented bot policy and a verified clean technical foundation.

03

Entity Identity and Schema

Weeks 3 to 7. Establish or correct the brand's entity presence in the Google Knowledge Graph and Wikidata. Implement complete Schema.org Organization markup with sameAs links to Wikidata, LinkedIn, Crunchbase, the brand's industry registries, and any other authoritative profiles. Map and implement Article, Person, FAQPage, Product and LocalBusiness schema as relevant. This is administrative work but produces lift across every AI engine simultaneously.

04

Author Authority and Editorial Standards

Weeks 4 to 9. Identify three to five hero authors. Build complete profile pages with linked credentials, ORCID where applicable, complete Person schema, and curated bodies of work. Publish an editorial standards page covering the role of AI, the review process, the credentials of the editorial team, and the disclosure conventions. Update existing key content with named bylines.

05

Distributed Presence Programme

Weeks 6 to 12. Audit and build presence on the platforms LLMs already trust. Wikidata first. Wikipedia where notability supports it. LinkedIn for hero authors and the brand. YouTube where the audience uses video. G2 where applicable. Reddit through credible long-form participation, with a contribution-to-promotion ratio of at least nineteen to one. This step accumulates over many quarters; the first ninety days establish the foundation.

06

Continuous Monitoring and Iteration

Week 10 onwards. Establish a quarterly cadence for citation-share review, prompt expansion, content updates, and technical re-audits. Identify the prompts where the brand should be cited but is not, and treat each as a research and content programme. The discipline is to extend the baseline established in step one, not to produce parallel reports for each new tool that arrives in the market.

How to sequence with limited resources

For brands with limited capacity, the priority order is clear. Step one is non-negotiable; without measurement, no other work can be evaluated. Step two is the most underrated single source of immediate lift, because it remediates problems the brand often did not realise it had. Step three is the most valuable administrative investment available. Step four produces the trust signals AI surfaces increasingly require. Steps five and six can begin in parallel with the others but reach their compounding benefit over multiple quarters.

What to expect in the first six months

Realistic expectations matter. From engagement data, a competently executed ninety-day programme typically produces measurable AI citation lift within four to six months, with the largest single contributor being the schema and entity work in step three. Authority and distributed-presence work in steps four and five tends to compound more slowly, with the visible benefit landing in months six through twelve and continuing thereafter.

The single biggest risk to programme success is impatience. The brands that see compounding returns are the ones that maintain the discipline through the first two quarters when results are still partial. The brands that abandon programmes after twelve weeks because the numbers have not yet moved are giving up on a capital investment one third of the way through its useful life.

Patience, in this discipline, is now an underrated competitive advantage.

ENGAGEMENT PRINCIPLE

Strategic Checklist.

Two checklists. The first is for the leadership conversation. The second is for the team that will implement the work. Either can be used independently. Used together they form a complete brief.

For the boardroom

The questions your board, executive team, or investors should be asking, with the answer that a credible AI visibility programme can produce.

-
- ☐ **Do we have a measurement baseline for AI search visibility?**
If not, the next ninety days should establish one. Without it, every subsequent investment is unaccountable.
 - ☐ **Are we accidentally blocking AI search bots?**
An audit takes two weeks. The visibility lift, if blocks are found and addressed, is often material.
 - ☐ **Does our brand exist as a coherent entity in the Knowledge Graph and Wikidata?**
If not, this is the most valuable single administrative investment available.
 - ☐ **Do we have credible named authors with documented expertise?**
Author authority is now consumed by AI surfaces as a discrete trust signal.
 - ☐ **Have we audited our presence on Reddit, Wikipedia, YouTube, and LinkedIn?**
These platforms account for the majority of AI citations in many categories.
 - ☐ **What is our editorial governance for AI-assisted content?**
The standard the discipline now requires is demonstrable accountability, with a public editorial standards page and named reviewer responsibility.
 - ☐ **How are we measuring AI citation share against our competitors?**
Tooling matured in 2025. Mid-market budgets now access this data.
 - ☐ **What is our planning horizon?**
The compounding period for authority is two to three years. Plans on shorter horizons mistake activity for progress.
-

For the practitioner

The execution items, sequenced by dependency. This list is intended as a working document. Use, mark, refine.

- ☐ **Connect Bing Webmaster Tools AI Performance dashboard.**
The most actionable single source of public AI search analytics.
- ☐ **Verify Google Search Console AI Overview impression data is collecting.**
Limited but improving signal.
- ☐ **Establish a baseline citation-share measurement across the top one hundred prompts in your category.**
Use Profound, Otterly, Peec AI, AthenaHQ, Scrunch, or BrightEdge AI Catalyst as appropriate.
- ☐ **Audit robots.txt with explicit per-bot decisions.**
Differentiate training, search, and user-fetch bots. Document the policy.
- ☐ **Audit WAF and CDN rule sets for hidden AI bot blocks.**
Cloudflare, Akamai, Imperva managed rule sets are the most common silent culprits.
- ☐ **Create or correct your Wikidata record.**
Lower notability bar than Wikipedia. More valuable than most teams realise.
- ☐ **Implement complete Schema.org Organization markup with sameAs links.**
Connect to Wikidata, LinkedIn, Crunchbase, industry registries.
- ☐ **Implement Article schema with author, datePublished, dateModified, citation.**
Standard for editorial content.
- ☐ **Implement Person schema for hero authors with linked credentials.**
Three to five hero authors per brand.
- ☐ **Publish an editorial standards page at a stable URL.**
Cover AI use, review process, editorial team, disclosure conventions.
- ☐ **Update key cornerstone content with named bylines and explicit dates.**
Visible publication and revision dates are freshness signals.
- ☐ **Front-load the answer in cornerstone pages.**
Forty-four per cent of citations come from the first third of cited pages.

- ☐ **Add primary citation, statistic, and authoritative quotation per major section.**
The Princeton GEO study lift is between thirty and forty per cent for each.

- ☐ **Audit and complete YouTube transcripts and metadata.**
YouTube is the single most-cited domain in Google AI Overviews.

- ☐ **Make PDF assets text-extractable with metadata and dates.**
Bing AI grounding actively reads PDFs. Image-only PDFs are invisible.

- ☐ **Establish a credible Reddit presence in relevant subreddits.**
Contribution-to-promotion ratio of nineteen to one minimum.

- ☐ **Build LinkedIn footprint for hero authors with professional positioning.**
LinkedIn is consumed across every major AI surface.

- ☐ **For Australian regulated industries: audit registry footprint.**
AHPRA, ASIC, TPB and similar registries carry weight in AI surfaces.

- ☐ **Set quarterly review cadence for citation share, content updates, and bot policy re-audit.**
The discipline is continuous, not project-bounded.

Key Takeaways.

If only one page of this report leaves the office on Monday, this is the page we would choose.

8%

Click rate on Google searches with an AI summary, versus 15 per cent without.

PEW RESEARCH CENTRE, JULY 2025

+40%

Average lift in AI citation visibility from adding statistics, citations, and quotations.

AGGARWAL ET AL., ACM SIGKDD KDD 2024

54.5%

Share of Google AIO citations now sourced from the organic top ten ranking pages.

BRIGHTEDGE GENERATIVE PARSER, SEPTEMBER 2025

800M

Estimated weekly active users of ChatGPT by October 2025, with Gemini at approximately 450 million MAU.

OPENAI DISCLOSURES & SENSOR TOWER

11%

Domain overlap between top-cited sources in ChatGPT and Perplexity, across 680 million citations.

PROFOUND, AUGUST 2024 TO JUNE 2025

3 to 13%

Range of citation URLs hallucinated by commercial LLMs, across cross-model audits.

PENN ET AL., ARXIV 2603.03299, 2026

Ten things worth committing to memory

1. Search did not break. The interface fragmented, and the criteria for visibility consolidated.
2. The same E-E-A-T signals do double duty in organic and AI surfaces. The split between SEO and GEO is largely artificial.
3. Each AI engine cites differently. Multi-engine measurement is now the baseline.
4. The cheapest visibility increase available to most brands is presence on the platforms LLMs already trust. Wikipedia, Wikidata, Reddit, YouTube, LinkedIn, G2.
5. The Australian rollout lag is a window. It is closing.
6. Most brands are blocking the AI bots they should be allowing. Audit before content investment.
7. The AI content gold rush is over. The expert content era has begun.
8. AI traffic is high intent but small in absolute share. Treat it as additive, not as a reason to reallocate budget.
9. Agentic browsing is real, public, and benchmarked. Reliability still trails human performance, but the direction is clear.
10. The brands that win the next decade of search will not be the ones that move first. They will be the ones that build authority slowly enough to be defensible.

Full Source Library.

Every primary document, peer-reviewed paper, and industry study cited in this report. Sources are presented in topical groupings rather than strict alphabetical order, to help the reader navigate by subject. Every URL was live at the time of publication. Where data is contested or methodology is partial, this is noted in the body of the report rather than in this reference list.

01 Google Search Central. What web creators should know about our March 2024 core update and new spam policies.
<https://developers.google.com/search/blog/2024/03/core-update-spam-policies>

02 Google Search Central. Updating our site reputation abuse policy.
<https://developers.google.com/search/blog/2024/11/site-reputation-abuse>

03 Google. AI Mode in Google Search: Updates from Google I/O 2025.
<https://blog.google/products-and-platforms/products/search/google-search-ai-mode-update/>

04 Google. Introducing AI Overviews in Australia.
<https://blog.google/intl/en-au/company-news/outreach-initiatives/ai-overviews-australia/>

05 Google APAC. Google AI Mode for marketers in Australia.
<https://business.google.com/aunz/think/ai-excellence/google-ai-mode-australia-marketers/>

06 Google Search Central. Search Quality Rater Guidelines (Sept 2025 revision).
<https://services.google.com/fh/files/misc/hsw-sqrg.pdf>

07 Google. Our latest update to the quality rater guidelines: E-A-T gets an extra E.
<https://developers.google.com/search/blog/2022/12/google-raters-guidelines-e-e-a-t>

08 OpenAI. Introducing ChatGPT search.
<https://openai.com/index/introducing-chatgpt-search/>

09 OpenAI. Introducing Operator.
<https://openai.com/index/introducing-operator/>

10 Anthropic. Introducing Contextual Retrieval.
<https://www.anthropic.com/news/contextual-retrieval>

11 Anthropic. Introducing Claude 4.
<https://www.anthropic.com/news/claude-4>

12 Claude. Introducing web search on the Anthropic API.
<https://claude.com/blog/web-search-api>

13 Microsoft Bing. Introducing AI Performance in Bing Webmaster Tools Public Preview.
<https://blogs.bing.com/webmaster/February-2026/Introducing-AI-Performance-in-Bing-Webmaster-Tools-Public-Preview>

14 Microsoft Copilot Blog. Bringing the best of AI search to Copilot.
<https://www.microsoft.com/en-us/microsoft-copilot/blog/2025/11/07/bringing-the-best-of-ai-search-to-copilot/>

15 llmstxt.org. The llms.txt proposal (Jeremy Howard, Answer.AI).
<https://llmstxt.org/>

16 Brave. Brave Leo AI overview.
<https://brave.com/leo/>

17 Perplexity. What is Pro Search.
<https://www.perplexity.ai/help-center/en/articles/10352903-what-is-pro-search>

18 Aggarwal et al.. GEO: Generative Engine Optimization (ACM SIGKDD KDD 2024).
<https://arxiv.org/abs/2311.09735>

19 Han et al.. Retrieval-Augmented Generation with Graphs (GraphRAG).
<https://arxiv.org/abs/2501.00309>

20 arXiv. RAG vs. GraphRAG: A Systematic Evaluation and Key Insights.
<https://arxiv.org/html/2502.11371v3>

21 arXiv. A Systematic Review of Key Retrieval-Augmented Generation Systems.
<https://arxiv.org/pdf/2507.18910>

22 Zhang et al.. Comparative Analysis of Faithfulness Metrics and Humans in Citation Evaluation (LLM4Eval@SIGIR 2024).
<https://arxiv.org/pdf/2408.12398>

23 Penn et al.. Cross-Model Audit of Reference Fabrication in AI-Assisted Academic Writing.
<https://arxiv.org/pdf/2603.03299>

24 arXiv. Detecting and Correcting Reference Hallucinations in Commercial LLMs and Deep Research Agents.
<https://arxiv.org/pdf/2604.03173>

25 Pew Research Centre. Do people click on links in Google AI summaries? (n=900, March 2025).
<https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>

26 SparkToro & Datos. 2024 Zero-Click Search Study.
<https://sparktoro.com/blog/2024-zero-click-search-study-for-every-1000-us-google-searches-only-374-clicks-go-to-the-open-web-in-the-eu-its-360/>

27 Reuters Institute. Digital News Report 2025.
<https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025>

28 Reuters Institute. Generative AI and News Report 2025.
<https://reutersinstitute.politics.ox.ac.uk/generative-ai-and-news-report-2025-how-people-think-about-ais-role-journalism-and-society>

29 Edelman. Trust Barometer Flash Poll: Trust and AI at a Crossroads, 2025.
<https://www.edelman.com/insights/familiarity-breeds-trust-in-ai>

30 Gartner. Gartner Predicts Search Engine Volume Will Drop 25% by 2026 (analyst forecast).
<https://www.gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents>

31 Ahrefs. Top 10 Most Mentioned Domains for ChatGPT, Perplexity and AI Overviews (78.6M queries).
<https://ahrefs.com/blog/top-10-most-cited-domains-ai-assistants/>

32 Semrush. The Most-Cited Domains in AI: A 3-Month Study (230,000+ prompts).
<https://www.semrush.com/blog/most-cited-domains-ai/>

33 BrightEdge. AI Overview Citations Now 54% from Organic Rankings.
<https://www.brightedge.com/resources/weekly-ai-search-insights/rank-overlap-after-16-months-of-aio>

34 BrightEdge. Ultimate Guide to AI Overviews.
<https://www.brightedge.com/ai-overviews>

35 Profound. AI Platform Citation Patterns (680 million citations).
<https://www.tryprofound.com/blog/ai-platform-citation-patterns>

36 Lantern. 10 Most Cited Domains across ChatGPT, Perplexity, Gemini and Claude.
<https://www.asklantern.com/blogs/10-most-cited-domains-across-chatgpt-perplexity-gemini-and-claudee-here-s-the-pattern>

37 SE Ranking. Google's AI Overviews: Updates and changes.
<https://seranking.com/blog/ai-overviews/>

38 Cloudflare. From Googlebot to GPTBot: who's crawling your site in 2025.
<https://blog.cloudflare.com/from-googlebot-to-gptbot-whos-crawling-your-site-in-2025/>

39 Similarweb. AI Discovery Surges: 2025 Generative AI Report.
<https://www.similarweb.com/blog/marketing/geo/gen-ai-stats/>

40 Digiday. In Graphic Detail: The state of AI referral traffic in 2025.
<https://digiday.com/media/in-graphic-detail-the-state-of-ai-referral-traffic-in-2025/>

41 ConvertMate. GEO Benchmark Study 2026.
<https://www.convertmate.io/research/geo-benchmark-2026>

42 ALM Corp. ChatGPT Citations: 44% Come From the First Third of Content.
<https://almcorp.com/blog/chatgpt-citations-study-44-percent-first-third-content/>

43 ALM Corp. Bing Webmaster Tools AI Performance Report Guide.
<https://almcorp.com/blog/bing-ai-performance-webmaster-tools-complete-guide/>

44 HouseFresh. How Google decimated HouseFresh.
<https://housefresh.com/how-google-decimated-housefresh/>

45 PPC Land. HouseFresh achieves notable traffic recovery after Google algorithm impacts.
<https://ppc.land/housefresh-achieves-notable-traffic-recovery-after-google-algorithm-impacts-2/>

46 PPC Land. Google disputes Pew study showing AI Overviews reduce clicks.

<https://ppc.land/google-disputes-pew-study-showing-ai-overviews-reduce-clicks-by-half/>

47 Search Engine Land. Google's AI Overviews are hurting clicks: Pew study.

<https://searchengineland.com/google-ai-overviews-hurting-clicks-study-459434>

48 Search Engine Journal. Bing Adds GEO To Official Guidelines, Expands AI Abuse Definitions.

<https://www.searchenginejournal.com/bing-adds-geo-to-official-guidelines-expands-ai-abuse-definitions/568442/>

49 Columbia Journalism Review. Reddit Is Winning the AI Game.

<https://www.cjr.org/analysis/reddit-winning-ai-licensing-deals-openai-google-gemini-answers-rsl.php>

50 MIT Technology Review. OpenAI launches Operator.

<https://www.technologyreview.com/2025/01/23/1110484/openai-launches-operator-an-agent-that-can-use-a-computer-for-you/>

51 Engadget. Atlassian is buying Arc maker The Browser Company for \$610 million.

<https://www.engadget.com/big-tech/atlassian-is-buying-arc-maker-the-browser-company-for-610-million-145236236.html>

52 Information Age (ACS). Google rolls out genAI search results in Australia.

<https://ia.acs.org.au/article/2024/google-rolls-out-genai-search-results-in-australia.html>

53 Connectedm. Google Rolls Out AI Mode Search in Australia (8 October 2025).

<https://www.connectedm.com.au/post/google-ai-mode-in-australia>

54 ADM+S Centre (ARC Centre of Excellence). AI Overviews have transformed Google search.

<https://www.admscentre.org.au/ai-overviews-have-transformed-google-search/>

55 Box Blog. Contextual Retrieval in Retrieval-Augmented Generation.

<https://blog.box.com/contextual-retrieval-in-retrieval-augmented-generation-rag>

56 PE Collective. Best Embedding Models 2026: MTEB Benchmarks & Pricing.

<https://pecollective.com/tools/best-embedding-models/>

57 Schema App. The Semantic Value of Schema Markup in 2025.

<https://www.schemaapp.com/schema-markup/the-semantic-value-of-schema-markup-in-2025/>

58 TheNextWeb. Google Cloud Next 2026: AI agents, A2A protocol, Workspace Studio.

<https://thenextweb.com/news/google-cloud-next-ai-agents-agentic-era>

59 Anthropic API Docs. Claude Platform release notes (citations API, web search, computer use).

<https://platform.claude.com/docs/en/release-notes/overview>

60 Yahoo Finance. ChatGPT user growth has slowed (Sensor Tower).

<https://finance.yahoo.com/news/chatgpt-user-growth-slowed-report-200603601.html>

61 Marketplace. ChatGPT now has 400 million weekly users.

<https://www.marketplace.org/story/2025/02/20/chatgpt-now-has-400-million-weekly-users-and-a-lot-of-competition>

62 TechCrunch. Reddit says it's made \$203M so far licensing its data.

<https://techcrunch.com/2024/02/22/reddit-says-its-made-203m-so-far-licensing-its-data/>

Citations within the body of this report follow inline academic conventions. Where a claim is supported by multiple sources, the strongest is cited inline and the corroborating sources appear in the dossier underlying this report. The full claim-validation register, including statistics flagged as contested or partially confirmed, is available on request.

About SKOMA Digital.

SKOMA Digital is an Australian advisory practice working with mid-market and enterprise brands on the practical work of being found, cited, and trusted in AI-mediated search.

The practice was founded by Sami Khan and is based in Australia, working with clients locally and internationally. The discipline of the work is straightforward. We help brands build the entity, content, and technical foundations that produce durable visibility across the search surfaces that exist today, and a defensible position for the surfaces that are still emerging.

The practice declines engagements where the brief is to produce content at scale without expert review. The practice declines engagements where the metrics are misaligned with what the data shows actually drives visibility. The practice does not promise rapid results. The practice does promise honest measurement, transparent reasoning, and work that compounds.

SKOMA's editorial standards, including the role of AI in research and drafting, the review process, and disclosure conventions, are documented publicly at skomadigital.com.au/editorial-standards

How this report was produced

This report is the synthesis of extensive research across primary documentation, peer-reviewed papers, industry studies, and SKOMA's own engagement data. AI assistants were used in the research and drafting stages, in the manner described in Chapter Twelve. Every claim of fact was reviewed against its primary source. Every URL in the source library was verified live at the time of publication. The editorial process is the same one we recommend to clients.

How to engage with us

If this report changes how you think about AI search visibility, the natural next step is a structured baseline assessment of where your brand currently stands across the five layers of the SKOMA Visibility Framework. The assessment takes two to three weeks, costs less than the average month of unfocused content production, and produces a written report with a sequenced action plan. To begin, write to sami@skoma.com.au or visit skomadigital.com.au. The work is durable. The conversations are honest. We would welcome the opportunity to do the work with you.

THANK YOU FOR READING.



Search did not break. It fragmented. The work is to build authority that holds across surfaces.

SKOMA Digital is an Australian advisory practice working with mid-market and enterprise brands on the practical discipline of being found, cited, and trusted in AI-mediated search. We would welcome the chance to do this work with you.

EMAIL

SAMI@SKOMA.COM.AU

WEB

SKOMADIGITAL.COM.AU

EDITION

VOLUME ONE, MAY 2026