

ARTEFACT

How is AI changing SEARCH?

What it means for
customers, marketers
and brands.

By Charlie Kay
Head of SEO and DEI Lead





25
COUNTRIES

1700
EMPLOYEES

+1000
CLIENTS

Artefact is a global leader in consulting services, specialized in data transformation and data & digital marketing, from strategy to the deployment of AI solutions. We are offering a unique combination of innovation (Art) and data science (Fact).



STRATEGY & TRANSFORMATION | AI ACCELERATION | DATA FOUNDATIONS & BI
IT & DATA PLATFORMS | MARKETING DATA & DIGITAL

TABLE OF CONTENTS

It's all clickable so navigate to wherever you please!

01

Introduction - What Is Happening To Search?!

3

02

A Timeline Of Search Innovation - AI And Beyond

4

[Early Foundations \(1998–2015\): The Matching And Ranking Era](#)

[The Neural Leap \(2017–2019\): The One Where Search Gets A Brain](#)

[Platform Diversification \(2010–Now\): Search Spreads Beyond Search Engines](#)

[The Generative Shift \(2022–2023\): Search Starts To Talk Back](#)

[The AI-First SERP \(2024–2025\): Search Becomes A Workspace](#)

03

How Large Language Models Work And Their Impact On Search

13

04

From Ranking to Reasoning In Google: How AI Overviews And AI Mode Actually Work

16

[AI Overviews: Composed Answers At Scale](#)

[AI Mode: A New System Of Reasoning](#)

[AI Mode vs. AI Overviews: A Side-by-Side Comparison](#)

04a

Who Is Using AI Search (And Who Isn't): Why Adoption Matters

20

[AI search Traffic Analysis: seoClarity Research Study](#)

[Search Usage In 2025: The Datos x SparkToro Report](#)

[When People Are Using AI Assistants, Which Are Being Used Most?](#)

[Who Is Actually Using ChatGPT?](#)

[What Do All These Stats Mean?](#)

[The Retrieval Layer – Why Rankings Still Shape AI Answers](#)

[seoClarity Original Research Insight](#)

05

What Does AI In Search Mean For Marketers And Brands

32

[5.1 Deep Audience Research & Discovery – It's All In The Data](#)

[5.2 AI Visibility & Performance Measurement](#)

[5.3 Tailoring Your Strategy To The Platform – Because Not All AI Search Is Built The Same](#)

[5.4 Brand Building & Digital PR – Fuel Awareness, Mentions & Authority In AI Search](#)

[5.5 Reputation Management – Shaping Brand Perception In The Age Of AI](#)

[5.6 Trust Signals – Earning Your Place In The AI Ecosystem](#)

[5.7 Semantic Optimisation and Structured Data – Optimising Structure For Machines](#)

[5.8 Multi-Modal and Multi-Platform Optimisation – Make All Your Content Accessible](#)

[5.9 UX & Content Architecture – Designing For AI Discoverability At Scale](#)

[5.10 Crawlability & AI Agent Access – Don't Block The Bots That Matter](#)

[5.11 Feedback Loops & Model Monitoring – Track, Test, Learn Across AI Ecosystems](#)

[5.12 Agentic Interaction Readiness – Optimise for Decision-Making \(Without Clicks\)](#)

1. Introduction – What is happening to search?!

There's a lot of noise about how AI is going to change marketing – for agencies, consumers, and brands. Even working in SEO, and navigating updates day-to-day, I find the amount of content being shared on this topic overwhelming. This piece is my attempt to distil that noise into something cohesive. A piece of research to help people understand what's really happening, how we got here – and, most importantly, what to do about it.

I don't believe we're witnessing the death of SEO. We're at a point of major change and entering another chapter in the evolution of search. The acronym might change. Our focus might shift. We will likely enhance what we are doing and learn new skills. But the foundations of what we are trying to help brands do and achieve remain intact, with our work still representing a core pillar of digital marketing.

As has always been the case, the brands and businesses that will win are those who have the best products, do the best work, understand their customers, can adapt to a changing world, are good at storytelling, and provide genuine value.

So grab a cup of tea (or two) and settle in. This isn't a short piece – but it's long for a reason. There's a lot to understand. And I hope this gives you a clear, grounded view of how we got here – and where we go next.



Charlie Kay
Head of SEO and DEI Lead

2. A Timeline of Search Innovation – AI and Beyond

Search is always evolving – from how it looks, the way it works, and the features made available, to the way people search and the platforms they turn to. We're undoubtedly at an inflection point but it's been coming for a *long time*. The moment we are in is part of a much longer timeline of change, shaped by new technologies, shifting behaviour, rising expectations, and growth.

So before talking about how to adapt to these changes, I thought it'd be helpful to look back. This section explores the big things that have happened in search that brought us here – not just in how search engines have changed, but also how *search behaviour* has moved beyond them.

◆ Early Foundations (1998–2015): The Matching and Ranking Era

Pay particular attention to the *italicised terms* – these foundational shifts in search are super relevant in the AI world we find ourselves in today.

- **1998 – Google search is launched**

Not the first, but the most successful of its kind. A box, a button, and a ranked list of links that would go on to shape how the world accesses information.

- **2001 – Google adds machine learning to improve spelling correction**

This is still true today. Type “restuarant” and Google will suggest “Did you mean: restaurant?” – In 2001, this functionality used early machine learning models to learn from common spelling errors and improve the interpretation of queries. Showing early signs of Google's vision to best understand a user's query and present the right answers.

These are results for *restaurant*
Search instead for *restuarant*

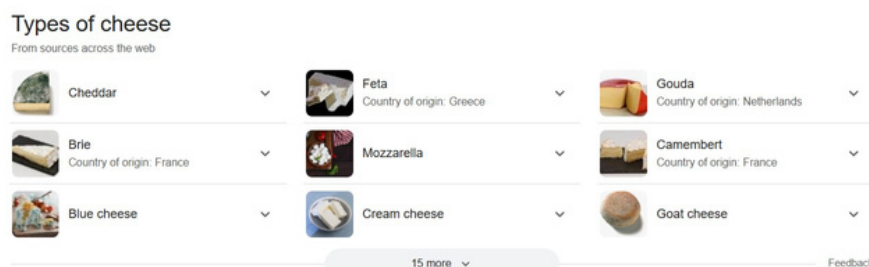
- **2003–2009 – Major algorithm updates: Florida, Panda, Penguin**

At the time, these were big updates. They targeted low-quality content, keyword stuffing, and suspect backlinks - marking a shift from rewarding more easy-to-manipulate optimisation tactics to *relevance-focused ranking*. This is where Google began prioritising the *quality of content* and the *intent of the user* to provide the most relevant results.

- **2012 – Google introduces the Knowledge Graph**

The Knowledge Graph is a massive database of entities - examples of entities in the knowledge graph include 'Book', 'Book Series', 'EducationalOrganization', 'Event' etc. At this point search shifted from *string-matching* to *entity recognition*. This meant search could better understand the 'real-world' meaning behind words – like people, places, and products – and represent them visually in rich SERP (search Engine Results Page) panels.

Google could now understand the real intent behind a query and surface content to align with the real-world concept or entity the user is asking about. [The image below shows information about cheese pulled from Google's Knowledge Graph in a SERP.]



● 2014 – Google introduces E-A-T (Expertise, Authoritativeness, Trustworthiness)

Introduced as part of Google's *Search Quality Evaluator Guidelines* (guidelines used by real human evaluators to rate the quality of search results). E-A-T emphasises the importance of *credible, accurate, and trustworthy* content – especially for topics like health, finance, and safety known as “Your Money or Your Life” (YMYL).



In the world of AI search, E-A-T principles are arguably even more important. LLMs like ChatGPT, and features like AI Overviews that use LLMs are trained to prioritise content that is factually accurate, cites credible sources, and demonstrates author trust signals – even if they're not specifically using the E-A-T framework.

● 2015 – Google launches RankBrain

This was Google's first significant use of machine learning in ranking. RankBrain helped Google better understand previously unseen queries by learning *patterns and relationships between concepts*. This was a major step towards real *semantic understanding*. Rank Brain was then able to match queries to the Knowledge Graph framework more effectively, returning even more accurate and rich results.

🧠 These foundational shifts laid the groundwork for what we now call semantic search – where engines interpret meaning, not just match words.

◆ The Neural Leap (2017–2019): The One Where search Gets a Brain

● 2017 – Google invents the Transformer

The ‘big brains’ at Google invented the Transformer - a ‘neural network architecture’ (a system which defines how data flows and is processed) that changed AI. It became the foundation for Large Language Models (LLMs) – AI that is trained on massive datasets to understand and generate human language. These LLMs have become the foundation of AI search today (more on this later).

This eventually paved the way for:

- **Google's BERT (Bidirectional Encoder Representations from Transformers)** – this helped search understand language in context – so within the full context of a sentence, rather than in isolation.
- **Open AI's GPT** – this enabled ChatGPT to generate human-like responses to prompts.
- **Google's Gemini** – a multimodal model that understands not just text, but also images, video, and code.

● 2018 – Google Lens and visual search launch

With the launch of Google Lens you could search using an image. You can now point your phone at a plant and ask, "What is this?" and you'll be provided with an answer by search. Or you can take a picture of a jumper you like, add a search "this but in green" and results would be returned for that exact jumper in a different colour (*try it if you haven't*).

● 2019 – Google launches BERT

Google began using LLMs to better understand natural language queries. For example, instead of simply matching the keywords in a search like "can I book a train for my mum from Canterbury to Bromley South", BERT helps Google understand that "for my mum" is essential to the intent – distinguishing it from someone booking for themselves or needing additional accessibility information.

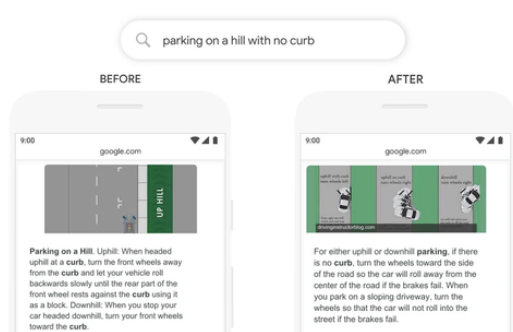
● 2021 – Google launches passage indexing

Passage ranking (also known as 'passage indexing') meant that Google would now rank individual passages from a web page even if it wasn't the main topic of a page if it best answers the user's query. This was a ranking, not an indexing, change impacting approximately 7% of all search queries.

● 2022 - E-A-T evolves into E-E-A-T (adding 'Experience')

Google updated its guidelines to include 'Experience' alongside Expertise, Authoritativeness, and Trustworthiness – reflecting the growing value placed on *first-hand perspectives, reviews, and user-generated content*. This shift aligned closely with the rise of Reddit, YouTube, and TikTok as trusted content formats in the SERPs. See next section for more on this.

💡 **BERT allowed Google to understand search queries better than ever before** – it was a breakthrough that paved the way for the conversational AI we see today.



"In the past, a query like this would confuse our systems—we placed too much importance on the word "curb" and ignored the word "no", not understanding how critical that word was to appropriately responding to this query. So we'd return results for parking on a hill with a curb!"

Google - *Understanding searches better than ever before*

◆ Platform Diversification (2010–Now): search Spreads Beyond search Engines

Parallel to the evolution of search Engines, user behaviour was also shifting – partially driven by the rise of new social platforms. Although not built as traditional search engines, they fulfil many of the same roles: *discovery*, *information gathering*, and *decision-making*. As user intent has expanded with the evolution of technology, so has the definition of what constitutes a ‘search’ experience. So not only has this meant that users, particularly the younger generation, have shifted platforms but also traditional search engines have responded by pulling more content from forums, video platforms, and social feeds directly into their results.

● 2010 onwards – Pinterest launches

Pinterest launches as a visual-first discovery engine, especially for interiors, fashion, DIY, and recipes. People tend to search in themes and by intent on Pinterest (e.g. “calm nursery ideas”), and content is ranked very differently – based on saves, engagement, and curation and not page rank or backlinks. More recently Pinterest have made several announcements about the use of AI on the platform – [introducing visual search](#) and enabling the creation of backgrounds but also [carefully labelling content created by Gen AI](#).

● 2010s onwards – YouTube becomes the go-to for product research and learning

For Millennials and Gen Z YouTube became (and still is in many cases) the first stop for how-tos, product demos, and, even life advice – especially for complex or visual topics. YouTube is the world’s second largest search engine, and as of early 2025 it boasts [2.53 billion monthly active users](#) and around [122 million daily active users](#). It is a really powerful way for brands to connect with audiences and build trust by showcasing the people, values, and expertise behind the brand.

● 2021 onwards – Reddit and forums rise in influence

Founded in 2005, Reddit is a community platform made up of smaller interest-based forums known as ‘subreddits’. Today, it sees over [365.4 million weekly active users](#). Since around 2021, search engines have increasingly surfaced forum content – with Reddit appearing more and more in search results. (In fact, [Google have a licensing agreement with Reddit](#) to crawl their content). This shift reflects a growing recognition by Google of the value in subjective, real-world questions like “is this worth it?” or “what’s it like to live in X?” – and a move toward prioritising lived experience over managed content.

● 2020 onwards – TikTok becomes a search engine for Gen Z

TikTok was launched in 2016, by 2020 had 700 million users, and is now estimated to have [1.8 billion monthly active users](#). It is a crucial place for brands to be, especially to connect with younger audiences because it’s where 40% of them now search! Whether it’s “what to do in Lisbon” or “best curling wands”, TikTok is increasingly used to discover content across all intents and purposes. TikTok ranks by recency, engagement, and authenticity.

🔍 Not only have these platforms reframed what search looks like and introduced new metrics by which to be ranked in traditional search but search has become more visual, social, and personal. From this point on, SEO optimisation was no longer just about ranking in Google, but about earning trust and attention where brand’s audiences actually are – forcing a shift in content strategy and the creation of roles specific to content for social media.

◆ The Generative Shift (2022–2023): Search Starts to Talk Back

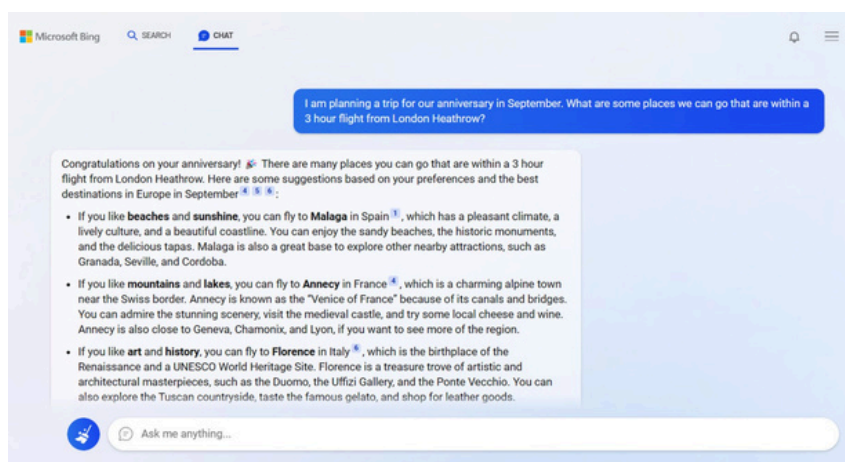
● Nov 2022 – OpenAI launches ChatGPT

ChatGPT is a large language model chatbot developed by OpenAI. It uses natural language processing to have human-like conversations and generate text. It was initially trained on GPT-3.5, OpenAI's large language model and generative pre-trained transformer. ChatGPT was not connected to the web at first, but was immediately popular, as demonstrated by its growth, and clearly demonstrated the power of Generative AI.

ChatGPT grew to **100 million users in two months** (that's faster growth than TikTok saw when it launched). Demonstrating the popularity and rate of adoption of Generative AI.

● Feb 2023 – Microsoft launches Bing Chat (later Copilot)

Bing was the first major search engine to embed OpenAI's tech into their experience by introducing a conversational assistant directly within its results. The chat was search-adjacent and could be used to expand on a 'normal' search. This move by Bing was a catalyst for the changing face of search -- forcing Google's hand to rapidly introduce changes to their search ecosystem.



● May 2023 – Google launches Bard (powered by LaMDA), and announces PaLM 2

Bard was Google's first attempt at a standalone conversational AI assistant and positioned as their answer to ChatGPT. Initially separate from search, it laid the foundation for integrating generative AI across Google's ecosystem. In mid-2023, Bard became available within search as an opt-in experience (via search Labs), offering AI-generated responses to queries alongside traditional results.

- *Bard helped Google test how users interact with AI in search – influencing the rollout of AI Overviews and Gemini integration.*

A wave of controversy followed the launch of Bard and Google's search Generative Experience (SGE), raising early concerns about the accuracy and transparency of AI-generated answers. One widely reported incident involved Google's own promotional tweet for Bard, in which the model incorrectly claimed that the James Webb Space Telescope took the first images of a planet outside our solar system — a mistake that led to a \$100 billion drop in Alphabet's market value.

This really highlighted the real risks of factual inaccuracy and the need for responsible deployment of AI in search.

Google also faced criticism for reusing full paragraphs of publisher content in its AI answers without consistently providing clear citations or links — fuelling debate around content ownership, visibility, and attribution in the age of AI search.

- **Dec 2023 – Google announces Gemini**

The successor to LaMDA and PaLM 2, Gemini is Google's new family of powerful multimodal LLMs, designed to reason across text, code, images, and more. It replaced Bard, and marked a significant step-change in performance across Google's AI ecosystem. More importantly, Gemini is what really prepared the ground for the next evolution of search — powering the shift from experimental AI answers to something much more native, integrated, and capable.

- **Late 2023 – Bing Deep search rolls out**

Bing describes Deep search as a tool to “provide even more relevant and comprehensive answers to the most complex search queries...an enhancement that offers the option for a deeper and richer exploration of the web... [It] takes the search query and expands it into a more comprehensive description of what an ideal set of results should include.” In many cases, for Bing, this is the point where *answers begin replacing links*.

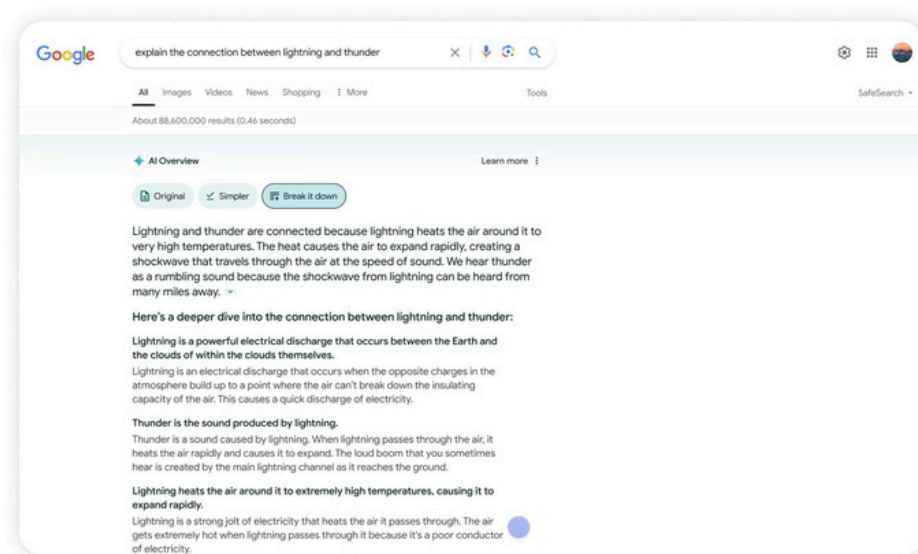
🧠 **You can think of all these things as search becoming interactive** — users stopped typing and scanning and started chatting, refining, and co-creating.

◆ **The AI-First SERP (2024–2025): Search Becomes a Workspace**

- **May 2024 – Google launches AI Overviews**

Powered by Gemini, Google started showing AI generated summaries above the standard search results — pulling data from multiple sources. AI Overviews include links, citations, and sometimes multi-step reasoning. As of now, the launch of AI Overviews marks the single biggest change ever made to Google search — replacing a list of links to third-party content with *generated answers*.

[The image at the top of the next page shows an example of an AI Overview in the Google SERP. The query is 'explain the difference between lightning and thunder'. The answer is compiled of 5 paragraphs -- an AI Generated answer in the SERP with no links.]



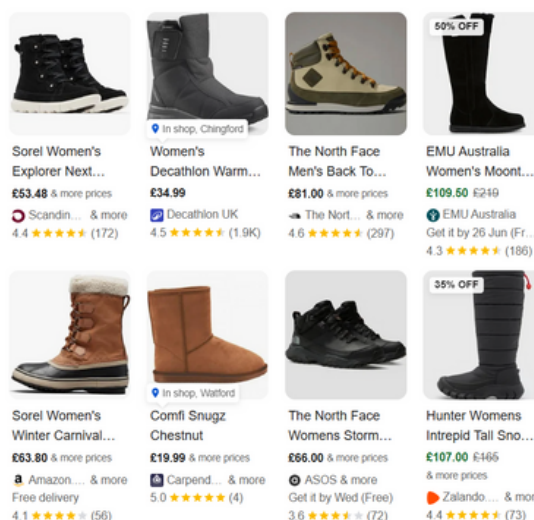
Late 2024 – AI-powered Shopping & Product Grids Expand

Search results become increasingly shoppable as AI curates dynamic product grids based on user intent, context, and past behaviour. These grids aren't just keyword-triggered – they're assembled by interpreting queries holistically (e.g. "eco-friendly trainers for winter running"), factoring in real-time availability, reviews, specs, and even individual preferences. AI now starts to act more like a personal shopper, filtering and ranking products based on *relevance, trust, and performance*.

- For retailers, this means structured data, product feeds, and high-quality product content (including UGC – User Generated Content – and reviews) are key
- Brands that lack product schema, missing prices or ratings, or unoptimised product pages risk exclusion from these surfaced results – even if they rank well organically.
- This change also expands across surfaces, the various platforms and interfaces within Google including classic search to Google Lens, YouTube Shopping, and AI Overviews etc. where product panels are dynamically generated.

Latest research shows that Organic Product Grids show for 14% of all queries on mobile and 15% on desktop in the US.

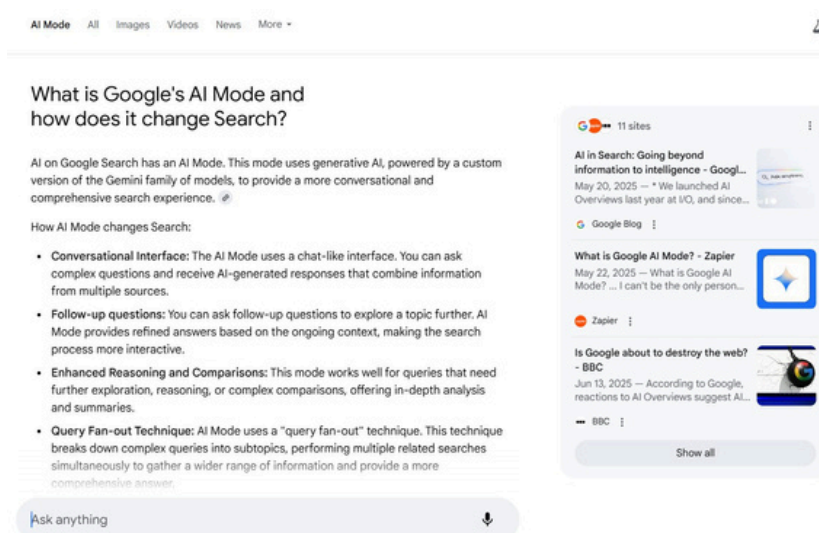
[The image to the right shows 8 products in an Organic Product Grid in the SERP for 'boots for winter in London in the UK']



May 2025 – Google I/O: Over 100 AI-related announcements, including:

- **AI Mode – An AI Assistant built into Google search as a separate tab**
Toggleable, interactive research experience. AI Mode has turned search into a dynamic dialogue – users can ask follow-ups, compare information, and get summarised insights from across the web, all within one interface. Think of it as ChatGPT built into Google, with access to real-time data, citations, and able to assist with a range of tasks like planning, coding, brainstorming etc.

[The image below shows a screenshot of Google's AI Mode released in May this year in the US. The prompt is 'What is Google's AI Mode and how does it change search? The answer below is made up of a paragraph and a bulleted list. To the right is a separate box with links to 11 different sites selected once the response to the query is built].



Search Live – Explore topics in real-time and across media types

Search Live allows users explore topics through spoken real-time conversations with search. It is multimodal, and contextual – it doesn't just show traditional text results, but also pulls from news, videos, social posts, podcasts, forums, and more. This means surfacing in search no longer depends on ranking a page – your brand's presence across formats and platforms now matters more than ever.

Gemini 2.5 – Google's most advanced, fastest LLM yet

Gemini 2.5, released in mid-2024, is *multimodal*, *memory-enabled*, and capable of deeper reasoning. It powers Google's Workspace features, AI Overviews, AI Mode and search enhancements with *better factual grounding and faster retrieval*. As it continues to integrate across products, Gemini reshapes how information is accessed – not just via links, but via *synthesised insights drawn from trusted sources*.

- **Agentic Checkout – AI can now guide and complete transactions.**

With Agentic Checkout, users can now ask AI assistants to help them research, compare, and *complete a purchase*, all within a *conversational interface*. For example, ChatGPT or Gemini can guide users from discovery (“Which hiking boots are best for wet terrain?”) to purchase (“Buy the top-rated waterproof pair in my size”).

🧠 **We’re now in the age of agentic search** – where AI doesn’t just help you search, it helps you decide, plan, and act.

- **Ask Photos – Use natural language to search your own Google Photos.**

With *Ask photos*, Google lets users search their photo library with queries like “Show me our beach trip in 2021 when Anna was wearing a red dress.” This showcases LLM-powered personal memory recall – combining visual understanding with contextual prompts.

← END **Summary: We’re Navigating a New search Ecosystem and We’re Optimising for AI**

So, as I hope this section has shown, AI has been shaping search for years – but the type, scale and speed of change we’re seeing now is different. It’s not just about new technology; it’s a redefinition of how, where, and why people search, and how search *operates*.

Today’s search ecosystem is no longer limited to Google or Bing. It spans TikTok, YouTube, Reddit, Pinterest, and a fast-growing wave of AI-powered assistants. Each of these platforms has its own discovery patterns, trust signals, and user expectations – and success now means understanding how to be visible, credible, and useful across all of them.

3. How Large Language Models (LLMs) Work, And Their Impact On search

We've looked back at how search has changed and evolved and what that has meant for marketing, users and brands. Armed with that knowledge – to understand how AI is truly changing search now – you first need to understand the engines powering them: large language models, or LLMs.

These models are the brains behind ChatGPT, Gemini, Claude, and every chatbot or AI-powered summary you've seen appear in search engines recently.

What is a Large Language Model (LLM)?

A Large Language Model (LLM) is a type of AI trained on huge amounts of text – books, websites, articles, forum posts, documentation, and more. They learn patterns in language from this documentation and use that knowledge to:

- Understand natural language queries,
- Generate fluent, human-like responses,
- Summarise or rephrase content,
- Infer relationships between words, concepts, and ideas.

When you're using ChatGPT it's easy to forget that LLMs don't actually 'know' things like a person, or 'read' content the way humans do. And, unlike traditional search engines, they don't query an index and return the results most relevant to a query. Instead, they generate responses based on patterns they've learned from training data.

Okay, okay. So How Do LLMs Actually Work?

It's important to understand how LLMs *actually* work to know how search is changing based on their introduction to it.

1. Tokenisation

Before training LLMs on large datasets, all of the training text has to be broken down into small units called 'tokens' – these can be whole words, parts of words, or even punctuation marks. These tokens are how the model sees the input – like translating language into machine-readable *chunks*.

For example:

"search is changing" might become → [search] [is] [chang] [##ing]

This ‘tokenisation’ helps the model process language mathematically. The model sees sequences of tokens and learns to predict what comes next. These tokens are then mapped to *vectors* (arrays of numbers) through something called an *embedding layer*. This allows the model to represent the *meaning* of words numerically—so it can learn, for example, that “puppy” and “dog” are more closely related than “puppy” and “car.”

2. Training on Huge Text Datasets

Once the datasets are tokenised the models can be trained on them. The core model is trained on *billions (sometimes trillions)* of tokens! Examples of the types of datasets that would be tokenised include publicly available or licensed content such as books, Wikis, news sites, documentation, forums like Reddit, and more.

These sources are referred to as the model’s training data. The model learns by predicting the next token in a sequence, millions of times — essentially completing sentences at scale.

This process helps the model learn:

- How grammar and sentence structure work,
- What facts are commonly repeated (e.g. “Paris is the capital of France”),
- How tone and style vary across contexts (e.g. news vs. blogs vs. conversations),
- How discourse flows logically (e.g. how to argue, list, explain),
- What concepts frequently appear together (e.g. “booking” and “hotel”),
- What questions people commonly ask — and how they’re answered.

3. Fine-Tuning

Once the model is trained, developers will often fine-tune it to improve the way it responds to real-world prompts. This is done in different ways:

- **Reinforcement Learning from Human Feedback (RLHF):** Where human reviewers rate model outputs for helpfulness and safety.
- **Instruction tuning:** The model is trained to follow user directions better to provide a better experience.
- **Custom fine-tuning:** Sometimes, models are tailored for a specific brand, tone, or purpose.

Well-known models that use fine-tuning, include **ChatGPT (GPT-3.5 / GPT-4)**, **Claude (Anthropic)**, **Gemini (Google)**, **Perplexity**, and **Bing Copilot** (based on GPT-4 Turbo) which are adapted for tone, safety, product integration, or use cases.

Fine-tuning improves safety, accuracy, and alignment with user needs — helping turn raw capability into useful, trustworthy interaction. (Although, as has been proven many times, accuracy is something all these models are still improving).

4. Generating a Response (Inference)

This is what happens when you ask a question or type a prompt into an AI Assistant or LLM-powered interface like ChatGPT, Claude, or Google's new AI Mode.

The model doesn't look up an answer from a database or perform a search (unless retrieval is enabled — see next section). Instead, it uses everything it learned during training to generate a response, one token (i.e. word fragment) at a time. This process is called 'inference'.

For example:

Prompt → "Best family hotels in Cornwall?"

Response → The model draws on patterns it's seen in similar queries to generate a plausible, fluent answer — not by retrieving, but by predicting what comes next, token by token.

5. Retrieval-Augmented Generation (RAG)

Most LLMs are trained on static datasets that end months or even years in the past. For example, when ChatGPT was first released, it was trained on data up until 2021. However, most well known models now incorporate retrieval-augmented generation (RAG). As a result this means they can access and incorporate live information while being used — blending learned knowledge with up-to-date sources.

How RAG works:

1. A query is inputted into the model,
2. It then performs a real-time web or document search,
3. It retrieves relevant sources (e.g. web pages, files, databases),
4. It uses this external information to generate a more detailed and accurate answer.

✓ **Models that use RAG today include:**

- **Perplexity** — aggressive real-time search + citation
- **Bing Copilot** — uses Bing index for retrieval (sometimes it might even use Google!)
- **Google AI Mode & Bard (Gemini)** — fetches real-time results
- **ChatGPT (with browsing or file upload)** — supports retrieval if enabled **Claude 3** — retrieval via uploaded documents and APIs

RAG improves factual accuracy, especially for live topics like news, pricing, or availability.

Why Does This Matter for Search

LLMs don't work like traditional search engines. They don't crawl websites and present a list of linked results — they generate the response themselves, based on what they've learned and (sometimes) what they can fetch in real time.

4. From *Ranking* to *Reasoning* in Google: How AI Overviews and AI Mode Actually Work

A lot of the conversation about AI search has been about what it looks like – a summary at the top of the search results, a chat box, fewer clicks, changing user behaviour, etc. But if we want to understand how to adapt, specifically to changes that are being seen from Google, we need to understand how these systems actually work.

Based on Google's recent release of new patents and very helpful analysis done by some of the industry's best (looking at you Michael King), we now know that **AI Overviews** and **AI Mode** are built and work differently to traditional Google search. Where SEO once focused on optimising pages for a query, with generative search we now have to optimise for a space that is hyper-personalised, intent-first, multi-query, and reasoning-driven.

Below is a simplified explanation of how these two systems operate, and what that means for visibility.

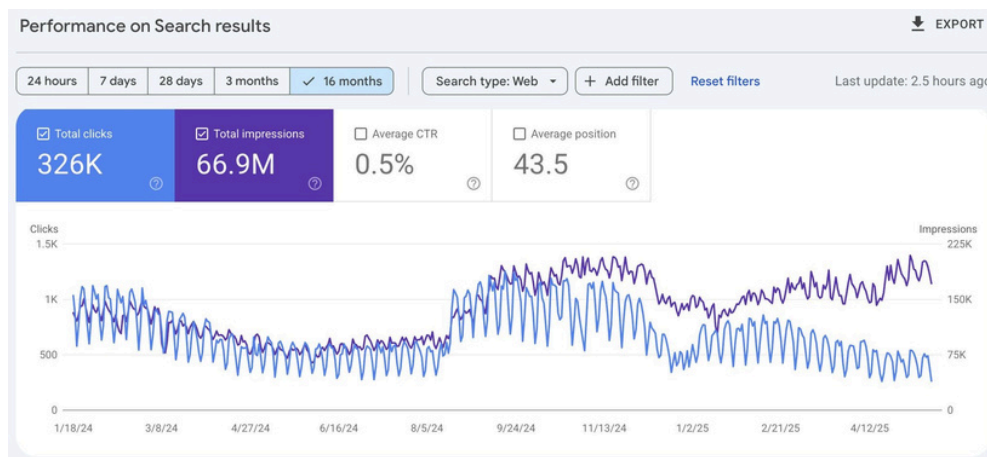
AI Overviews: Composed answers at scale

AI Overviews are everywhere now. According to [a study done by Ahrefs](#), they show for 9.46% of all keywords on Desktop and 54.61% of all searches by volume. They are triggered within the traditional search results and are mostly more 'extractive' than 'generative'. By that, I mean that Google is mostly surfacing reworded or repackaged information from real websites rather than synthesising information in the same way an AI Assistant might do:

1. Google starts by fetching documents for the original query – but also pulls in results for related, recent, and implied queries (this differs from traditional search where results would be pulled based on the understanding of the one query).
2. From this set of documents a summary is generated using an LLM.
3. Then citations are added if the summary can be verified against the retrieved documents (which is what we see most of the time with AI Overviews, links are added to sentences or paragraphs within the answer).

Alternatively, in a 'generative-first' approach – for example ChatGPT (without browsing) – the model produces a summary first and then after that searches for documents to verify or support what it says and then refines the output based on what it finds. This is a big change because if content isn't structured in a way the model can easily map to its own generated answer, it might be completely ignored.

AI Overviews have totally reshaped traditional search in terms of the way it looks, the way information is digested and in the impact it has on user behaviour. There have been so many examples of this shared (one shown below) - impressions are rising but clicks are falling. Users are getting the answers they need in the SERP and not clicking through to the source of the answer - websites.



[Image Caption: The *image* depicts a performance report from Google search Console that shows total clicks and total impressions over a 16 month period. The image clearly shows clicks declining and impressions rising].

In other cases, **both** impressions **and** clicks are falling. This is mostly seen when pages that previously ranked highly on SERPs are no longer cited. Previously having owned a featured snippet or ranking #1 doesn't guarantee visibility in AI Overviews because the systems work differently. Citation now depends on semantic clarity, factual accuracy, and how easily your content can be extracted and verified by the LLM.

AI Mode: A new system of reasoning

AI Mode is Google's full-screen, conversational experience. It's a chatbot and it's so far removed from the ranked list of links we call traditional search. It doesn't just find websites and content — it tries to understand what the user really wants, and then *generates an answer using multiple LLMs* in sequence. According to Google's patents, this involves:

- **Contextual interpretation:** Google factors in user history, device data, location, and account behaviour to enrich the query.
- **Synthetic query generation** (or "query fan-out"): Instead of searching for your original query, the system generates many related or implied versions to expand the intent space
- **Custom corpus retrieval:** Content is retrieved across these synthetic queries — not just by relevance to the original query, but to the full spread of possible intents.
- **Query classification:** The system decides if the user needs a comparison, an explanation, a decision etc. — and picks the most relevant LLM(s) accordingly.
- **Reasoning and response synthesis:** Responses are built paragraph-by-paragraph using reasoning chains — not just summaries but decisions about what to include, compare, and cite.

→ The model is actually trained to produce and evaluate its own reasoning steps — known as reasoning traces — before it settles on a final answer. That means content needs to help the model *think clearly*, not just answer cleanly.

- **Personalisation:** AI Mode also uses persistent user embeddings to personalise outputs at a deep level — this impacts not just what is shown, but how it's framed too.
- **'Pairwise' paragraph comparison:** A lesser-known step is that the model often compares your content directly against a competing paragraph to decide which better supports the user's intent. This isn't scoring — it's judgement. And it happens at the paragraph level.

AI Mode isn't the same as traditional search, *or at least that is what the patents say*. It doesn't return results — it builds them. Content isn't ranked, it's judged - interrogated, contextualised, reasoned through and generated. And this doesn't happen in isolation. It's evaluated for how well it contributes to a broader conversation.

LLMs might start with one article, move to another page, pivot through a table from another site, and end with a video clip from a fourth or more! This is no longer about ranking. It's about *being reasoned with*. AI Mode doesn't just retrieve and rank content — it engages with it critically and conversationally. This is a shift from transactional search ("find me X") to a more dynamic, interpretive process ("help me understand X through discussion and evidence") - more on this later.

Also, this isn't experimental anymore — this is already happening in production. AI Mode has been released into the wild in the US and will be expanded into other markets soon. In fact, at the time of writing this, Google has just announced the expansion of AI Mode to more users in the US. These models are live, serving millions of users.

AI Mode vs. AI Overviews: A Side-by-Side Comparison

Feature	AI Overviews	AI Mode
Trigger	Auto-triggered in search	User-initiated in Gemini/AI Mode
Interface	Embedded in classic SERPs	Full-screen, conversational UI
Query Expansion	Limited fan-out	Broad fan-out across multiple synthetic intents
Retrieval method	Standard index + some semantic scoring	Dense retrieval + paragraph level comparison
Reasoning model	Simple summarisation or verify-after-generate	Compositional reasoning with intermediate reasoning steps
Citation mechanism	Scroll-to-text or inline	Based on alignment with reasoning steps

Personalisation	Minimal (location, prior query)	Deep personalisation via persistent user embeddings
Output format	Static summary block with optional expansion	Dynamic responses: cards, timelines, explanations
Ranking logic	Extracted from ranked documents	Synthesised across sources – ranking is now selection
SEO implication	Content must be verifiable, structured, semantically rich	Content must be reasoning-ready, chunkable, and persona-aware

💡 Why It Matters

These systems do fundamentally change search and continue the evolution of it from:

Keywords → **intent clusters**

Ranking pages and paragraphs → **Synthesised paragraphs and answers**

Rankings → **reasoned inclusion**

CTR + Clicks → **citation and influence**

This isn't just a shift in layout, or an additional feature – it's a shift in logic. If you want visibility in the new landscape, your content has to work for *machines that synthesise*, not just index. And you're competing not just for attention, but for *relevance in a 'reasoning chain'*.

4a Who is Using AI search (And Who Isn't): Why Adoption Matters

It's important to look at how AI Assistants are actually being used, how much they are being used and by whom. Although there is no doubt that AI search adoption is growing it's not yet even *remotely* close to outpacing Google's 'traditional' search - in fact one of our partners seoClarity built a statistical model that triangulates searches on ChatGPT to about 2.5% of Google searches.

Not everyone is on board, and behavioural change takes time. After its launch in 1998 it took *eight years* for 'Google' to be added to the Oxford English Dictionary as a verb. I'm not saying it will take this long - we live in different times and change is happening fast, but there is still time to understand, prepare and take action to enter this.


So, before reacting to sensationalist headlines that are - ironically - designed for clicks, we need to look at the usage data. This is how brands and businesses can understand where opportunity lies and what this shift means for them. I've looked at several different data sets that show slightly different things to give a good view of where we are today.



AI Search Traffic Analysis: seoClarity research study

I'll be mentioning [seoClarity](#) quite a bit in this whitepaper, because they are awesome. And no, they haven't asked me to say this. I asked them if they'd like to be involved and they were, as always, incredibly generous with their time and data.


They recently completed a report on AI search Traffic analysis across 300+ domains in their Research Dataset. It shows that over *16.5 million visits* have been driven by the *top 4 AI search engines in under 9 months*, from September 2024 to April 2025. In the last 3 months alone traffic from AI search has grown by 127%. But, and this is an important but, less than 2% of all organic traffic is from AI search - with the help of Rand Fishkin, I'll dig a bit more into the share of market that AI search owns in the next section. [The chart below shows AI search Referral traffic from 300+ domains from September 2024 to April 2025. Growing from 75,321 referral sessions in September to 3,667,279 in April 2025.]




seoClarity then broke the data down by industry to see which receive the most traffic from AI search. As you can see in the chart below the industry breakdown is as follows, with information-rich sites receiving the most AI search traffic.

AI Search Traffic by Industry:

Education: 46.17%
Health: 14.42%
B2B Software: 12.14%
Finance: 9.34%
Ecommerce: 6.75%
Travel: 2.27%
Automotive: 1.13%



Finally, the report looks at the split of AI search traffic by AI search platform - ChatGPT and **Perplexity** lead in generating referral traffic; Gemini and Copilot lag behind. The chart below shows a bar chart by month from September 2024 to April 2025 with each bar split by 5 different AI search platforms ChatGPT, Perplexity, Gemini, Blackbox and CoPilot. It's clear from the chart that ChatGPT dominates the share of AI search referral traffic.



[The image to the left shows Share of AI Search Engines in Referral Traffic By Month. ChatGPT sends the most referral traffic by far, followed by Perplexity.ai]

💡 What Does This Mean?

While referral traffic from AI search is undoubtedly growing, it still only makes up a very small % share of total organic traffic. The amount of AI search referral traffic also varies by industry - being aware of AI search adoption for your specific industry is key.


Search Usage in 2025: The Datos x SparkToro Report

If you want another excellent piece of research on how search is being used in 2025, start with the [SparkToro x Datos State of search Q1 Report](#). The study is based on clickstream data from tens of millions of desktop users across the US, UK, and EU.


As Rand Fishkin puts it: *"It's filled with the graphs you'll be copy+pasting to your team, clients, and boss next time they overreact to a misleading headline."* (Little bit like how I've used them here, thanks Rand!)

Below are two charts based on the data from that report. They show that AI tool usage is steadily increasing across all regions — growing from 0.17% share in Jan 2024 to 0.55% share in May 2025 — but that traditional search still owns the majority share and AI still makes up a small slice of overall digital behaviour.

Search, AI, and E-Commerce platform visits in the US




Search, AI, and E-Commerce platform visits in the EU & UK



🖱 Clicks Still Happen (More Than You Think)

Another common misconception is that AI has wiped out the click. And while clicks *have* been declining, more data from the SparkToro x Datas report shows that for both the US and EU/UK, **clicking on organic results is still the most common outcome of a Google search.**



Look at the graphs again. Most people who search on Google **still click through to a third-party website from an organic result**. They *don't* click on a paid ad, they *don't* repeat the search, and they *don't* leave or bounce to another Google property. Similarly, as shown by data slightly further up, research conducted by seoClarity shows that AI search makes up less than 2% of total organic traffic (on average, this is higher for some industries).


I'm stressing this point because while zero-click searches are growing — and AI systems are clearly shifting behaviour — the majority of searches still result in someone visiting a third-party website.

This is *super* important context for SEOs and marketers: **AI isn't replacing search just yet, at the minute it's co-existing with it.**

🔍 AI Mode Specific Adoption: Still a very small part of overall Google usage

This point is echoed by clickstream data from SimilarWeb, [recently shared by Aleyda Solis](#). The analysis shows that **AI Mode usage remains small compared to other Google verticals**, such as Image, Video, News, or Shopping.

The takeaway is that AI Mode is growing, but still stabilising. It's not yet a dominant behaviour — it's an emerging one. The graph shows that over a period of 16 days since its launch AI Mode usage grew from 0.84% to 1.16% share.

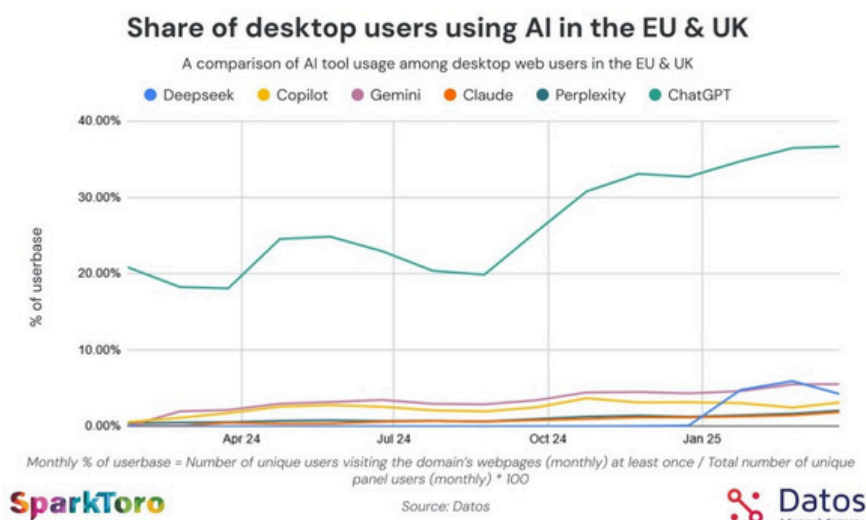


When People are Using AI Assistants, Which Are Being Used Most?

Despite the fact that AI platforms only make up a fraction of total digital usage, they are undoubtedly growing - and there are leaders in this race.

The SparkToro x Datos Report shows that ChatGPT has become the dominant AI assistant globally, making the Top 5 destinations visited from traditional search engines in both the U.S. and Europe - owning ~30% AI share.

[The chart shows several different AI tools' share of user base from March 2024 to March 2025 including Deepseek, Copilot, Gemini, Claude, Perplexity and ChatGPT. It clearly shows ChatGPT's growth from ~20% share of user base in 2024 to ~37% share of user base in March 2025].



This is supported by data from a report by CloudFlare '[Global expansion in Generative AI: a year of growth, newcomers, and attacks](#)' that shows ChatGPT is the number 1 ranked AI Service worldwide based on aggregated data from their 1.1.1.1 DNS resolver. The chart below shows the popularity rank of 11 different AI Services worldwide including ChatGPT, Character.AI, QuillBot, Codeium, GitHub Copilot, Deepseek, Perplexity, Claude/ Anthropic, Suno AI, Hugging Face and Grok.



💡 Why It Matters

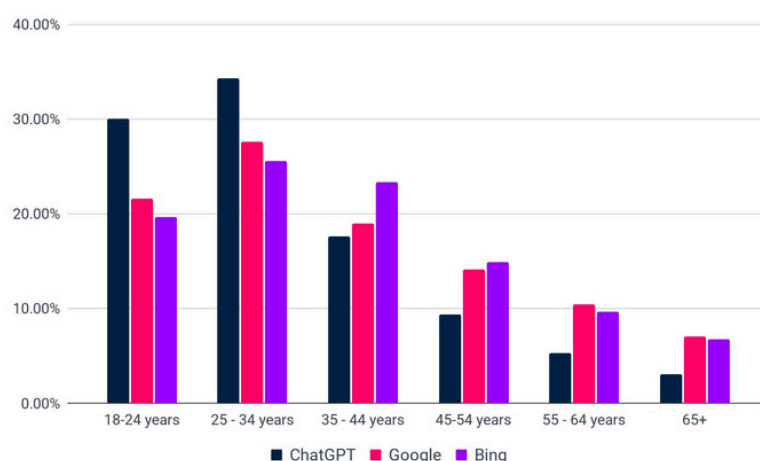
Knowing which AI Assistants and services are being used the most, how much they are being used in comparison to traditional search, and where they are being adopted is hugely valuable data. It allows us to understand how much resource should be dedicated to specific optimisation in these spaces. Research has also shown that different AI Assistants prioritise different signals when choosing which content to retrieve, repurpose and cite in responses.

👤 Who Is Actually Using ChatGPT?

While ChatGPT leads in overall usage, that usage is not evenly distributed across demographics or industries.

Age & Education:

- Adoption strongly correlates with higher education levels and digital literacy.
- The chart below shows the demographic split of usage for ChatGPT, Google and Bing with data taken from this [SimilarWeb report](#) and [this piece of research conducted by Ajprm](#).
 - It suggests that **63% of ChatGPT users are under 34 years of age**, meaning that Millennials and Gen Z are the biggest user groups, many of whom are daily users.
 - Only 13% of ChatGPT users are 55 or older and just 5% are 65+.




Gender:

Data from [Exploding topics](#) suggests that 64.3% of ChatGPT users are male, and 35.7% are female and [ChatGPT usage data from DemandSage](#) also suggests that usage is male orientated with a 54.7% to 45.3% split.

Location:

[Data updated June 15th 2025 by NerdyNav](#) shows that the 'United States leads in ChatGPT adoption with approximately 70 million monthly active users, representing about 19% of ChatGPT's global user base.




Occupation & Industry:

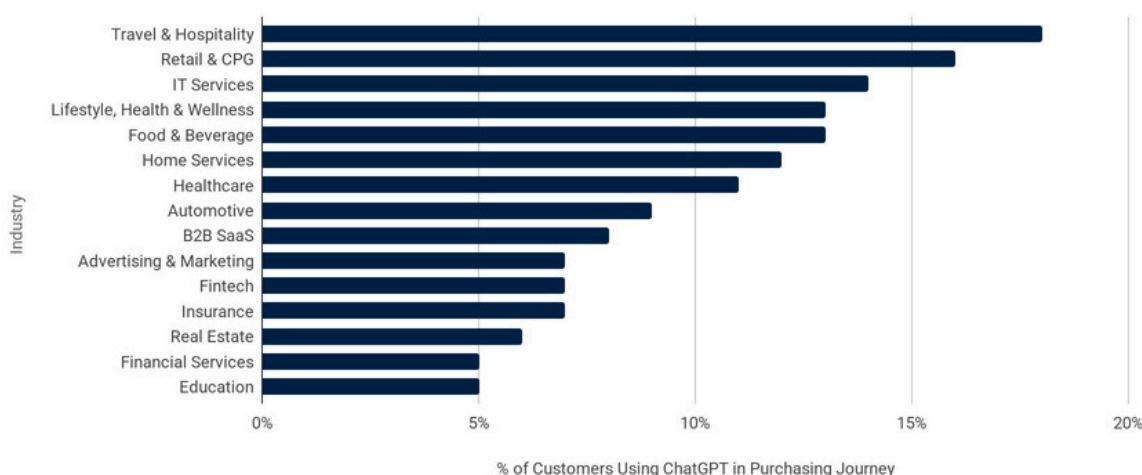
- 43% of U.S. adults say they've used ChatGPT for work.
- 92% of Fortune 500 companies have adopted ChatGPT in some capacity.
- Adoption is especially strong in marketing, design, tech, education, and finance.

Behaviour:

- The chart below shows data taken from a [First Page Sage ChatGPT Usage Statistics: June 2025 report](#). It shows that over the course of a year from May 2024 to May 2025 ChatGPT was mainly used for 'General Research'.




- Daily use is highest among knowledge workers and content creators who already use digital tools heavily.
- According to a study by SEMRush using real ChatGPT conversations, it tends to be used more for Programming and Social Media than anything else.
- More data from the [First Page Sage ChatGPT Usage Statistics: June 2025 report](#) shows % customers using ChatGPT in their purchasing journey by industry, with Travel and Hospitality and Retail and CPG leading the way at 18% and 16% respectively.



- And a study by Adobe of 5,000 U.S. consumers suggests that 39% have used Generative AI for online shopping with 53% planning to do so this year with referral traffic to Travel, Retail, and Banking websites specifically seeing substantial increases from Generative AI tools in Q1 2025.

The graph to the right shows referral traffic from AI search to three different industries Travel, Retail and Banking sent from ChatGPT from July 2024 to February 2025.



🧩 So What Do All These Stats Mean?

AI search is growing fast — but it hasn't yet overtaken traditional search, and it isn't used equally by everyone. Adoption is highest among younger, digitally fluent audiences, for commercial usage in specific types of industries like travel, retail, tech and IT, and for specific types of task. But it remains niche for others — especially older demographics, non-digital sectors, and companies without formal AI strategies.

⚠️ Barriers to Adoption – What's Slowing Down AI search Uptake among some audiences?

Despite the hype and rapid innovation, there are quite a few factors preventing widespread adoption — particularly in enterprise and mainstream user segments.

🌍 Concerns About The Environment

The power needed to train Generative AI uses an incredible amount of electricity, opening them up to public usage requires even more. In addition to electricity a substantial amount of water is needed to cool the physical computers used to run the AI.

🔒 Data Privacy & IP Concerns

41% of Americans feel AI does more harm than good when it comes to protecting personal information. Many organisations are wary of exposing proprietary data to LLMs, especially in sectors like legal, finance, and healthcare. There's still uncertainty around how user prompts and inputs may be stored, used, or trained on.

📰 Ethical Implications and Potential Misuse of AI


63% of adults are at least "somewhat worried" about the accuracy of AI-enhanced search engines and 68% fear misinformation; 62% are concerned about bias in AI search results.

🛡️ Regulatory and Compliance Risk

81% of people believe that AI needs stronger regulation. Concerns around GDPR, copyright, and AI regulation are rising.

This matters. It matters because brands and businesses should be marketing to their audiences, where they are and how they behave. Know your audience. It shapes where the opportunity is today, and how it will evolve. Brands shouldn't make knee-jerk decisions based on buzz statements, but instead should use data to inform their strategies and what they do. But I don't feel like I am telling you anything new here, this should have always been the case.

Bottom line: Adoption will be uneven — but, at some point, inevitable. It is predicted that AI search will grow 5x by 2030 and CAGR (Compound Annual Growth Rate) is predicted at 44.8%.



The graph to the left shows two data sets plotted in line. One for AI search and the other for Traditional search. The graph shows AI search growth outpacing traditional search growth by 2029.

🔄 Quickly Circling Back To RAG. The Retrieval Layer – Why Rankings Still Shape AI Answers

The data shows that, while AI search is growing it still doesn't boast *quite* the same usage rates as traditional search. Added to this, it's important to remember that what is powering many AI experiences is still **traditional search infrastructure**.

Language models like Gemini, ChatGPT, and Claude rely on Retrieval-Augmented Generation (RAG) — a process that I described above - where the AI search first expands your query (fan-out) and then pulls results from external knowledge sources before synthesising an answer.

In most consumer-facing AI search tools — especially Google's AI Overviews and AI Mode — that retrieval process leans heavily on **Google's core index and ranking systems** and ChatGPT on Bing's.

🧠 The AI doesn't just answer — it also retrieves information.


- AI systems retrieve from **traditional search indexes**, which means **ranking - and ranking well - in organic search still influences whether you're visible in AI-generated answers**.
- Gemini, Claude, and Perplexity often use similar logic: expand query → retrieve from web → compress into an answer

seoClarity Original Research Insight

A recent study conducted in May 2025 by [seoClarity](#) - our AI Tool partner who very kindly shared this data with us ahead of publishing - found that there is a “nearly linear relationship between traditional search ranking and AI Overview citations” – i.e. the higher your page ranks in classic search, the more likely it is to be cited by the AI system.

Key Takeaways:

- **97% of AI Overviews (AIOs)** include at least one URL from the Top 20 organic search results.
- **96% of AIOs** cite at least one page from the Top 10 organic rankings.
- On average, **5 web results overlap between AIO citations and Top 20 search results.**



The graph to the left shows the distribution of overlapping URLs in AI Overviews and traditional search.

- **50% average overlap** between AIO citations and Top 10 organic results.

The graph to the right shows the % overlap between AI Overviews and the Top 10 organic results. Showing there is a 50% average overlap.



- There is **direct correlation between ranking position and AIO inclusion**:
 - **Rank 1**: 54% likelihood of inclusion in AIO
 - **Rank 2**: 47%
 - **Rank 3**: 41%
 - **Down to Rank 20**: 7%

- **19% of keywords had 100% overlap** between AIO citations and Top 10 search results. **Industries with highest AIO overlap:**
 - **Healthcare** leads with the highest average citations (6.3 per keyword).
 - Followed by **Service, B2B, Finance, and Electronics**.

🔍 So what does this mean?

- You still need to *rank well* to be cited — even in the AI layer
- Your structured data, topical authority, and crawlability directly influence both ranking and retrievability

💡 *TL;DR: Don't neglect performance in traditional search. What LLMs know and where they get their data is still shaped by what they can retrieve — and they often retrieve what ranks.*

5. What Does AI in search Mean for Marketers and Brands — How to Adapt

We've looked at how search has evolved, how large language models work, how AI Overviews and AI Mode fundamentally reshape the search experience and the rates at which they are being adopted. So, we've reached the big (and heavily debated) question: **what should marketers, SEOs and brands do about it?**

While a lot has changed, the core principles of marketing remain the same — it's about being discoverable, useful, and trusted. But, in this new landscape, how you achieve that looks different.

The New Rules for Visibility in AI Search and Beyond

Search hasn't been limited to Google and Bing for some time — and with the rise of AI-powered assistants, the role of search engines is evolving even further. Today, it's less about being ranked first and more about being discoverable cross-channel and cited, summarised, and reasoned with by language models.

Ultimately, I think it's about creating stories and messaging that truly resonate with your audience and answering their questions and needs in the places they are, at the right time. If you have been doing SEO in silo and creating content solely for search - you have been doing it wrong.

5.1 Deep Audience Research & Discovery — It's All in the Data

Working at Artefact, we obsess over data — what it tells us, and more importantly, how we leverage it. That hasn't changed. What has changed, parallel to the evolution of search systems, is the type of data we now need to prioritise - particularly in the world of LLMs.

→ Know who you're taking to

Traditional search data is still valuable — there are enough people still using traditional search that it's representative of broader user intent and behavioural trends. And although keyword-level data has become less relevant as search has evolved, understanding topical demand and clustering is still invaluable in helping to guide content creation, especially when combined with more qualitative data from:

- CRMs
- Reviews
- Web or audience analytics software
- Forums (like Reddit)
- Comments on social media platforms like TikTok, YouTube etc. and
- Surveys.

Using this data provides insights into consumers' psyche - their opinions, thoughts, feelings, values and consequently the type of content that will be most valuable to them.

Why This Matters Even More in an AI Search World

As I touched on earlier, AI search doesn't treat all users the same. Many systems — from Google's AI Overviews to ChatGPT's personalised memory (if you have this feature enabled), adapt results so they are hyper-personalised - based on query context, user intent, search history, behavioural patterns etc. In other words, the same query may generate different outputs depending on who's asking. Or even generate different answers each time the same user asks the same question.

That's why deep audience understanding is no longer just a marketing advantage — it's an input signal to how and whether your content gets retrieved.

Key changes to consider

- **Intent diversification:** AI tools deconstruct complex queries into sub-intents (known as "query fan-out") and retrieve relevant content for each. As a result, if your content isn't comprehensive and doesn't cover multiple angles, it may be ignored. We will touch more on this later in "UX & Content Architecture".
- **Personalised query adaptation:** Prompts are often phrased conversationally or based on prior questions. That means you must understand how your audience asks, not just what they search.
- **The way that people search is changing** - this has been happening for a while — as demonstrated at the beginning of this whitepaper — but the advent of AI search has accelerated this change. Users are becoming more used to asking long questions, covering multiple angles, including personal details and having conversations. This is *very different* from existing search queries.
- **Data being used to inform AI search strategy** - We've been having really interesting conversations with seoClarity about this. If people aren't searching in the same way in AI search, how do we get data on what they are searching? seoClarity have built a way to measure demand for particular topics and 'queries' in AI search - this is what we need.
- **Location & familiarity bias:** LLMs can favour brands a user has interacted with before, or those prominent in their region or niche community.

What To Do

- **Expand traditional keyword research** to include **prompt analysis, intent segmentation**, query rewrites, data on questions that are asked, and look at the long tail questions people are asking. We are working closely with seoClarity and their exciting new tools to best understand how people are searching using AI Assistants and what demand looks like for different types of prompts, in different industries and for different types of intent.
- **Granularly segment your content strategy** to consider different personas, needs, entry points, intents and contexts — especially for top-funnel queries. This will also help when it comes to understanding what metrics to use to measure the performance of this content (more on this next).

- **Use forums, AI search, review sites, social media comment sections and tools like AnswerThePublic** to understand the language that your audience is using. We have to get granular, as traditional search Volume and tools like Keyword Planner are less relevant to determining how people interact in the digital world changes - using a broader research set will help.
- **Align content with regional colloquialisms, preferences in tone, and depth of information** to increase your visibility in spaces that may be used by AI search to generate an answer.
- Use performance data to analyse 'zero-click' or 'impression-only' mentions in AI Overviews and AI tools – these can indicate relevance even without direct traffic. Understanding the areas you are already considered relevant and the type of language used in these spaces can help with ongoing optimisation.

This leads us neatly into visibility and measurement in the world of AI search.

5.2 AI Visibility & Performance Measurement

The way that organic search success is measured has seen the same metrics used for years, and although I'm not going to jump on the 'Clicks are Dead' bandwagon, I will agree that with the growing influence of AI search, the standard reporting toolkit needs an upgrade. Not only is it critical to be able to demonstrate success across new platforms, and more generally for proper attribution, - without accurate data and measurement, you can't get the insight you need to continue improving strategy.

We've relied for years on metrics like Impressions, Clicks, Sessions, CTR, Average Position/Rank, Bounce Rate, Time on Page, Conversion Rate, number of Backlinks, Domain Authority etc. – but these won't accurately reflect performance in a world where users may never actually land on your site, and where their journey is increasingly shaped by answers generated elsewhere. We need new ways of understanding influence and performance. *So what does that look like?*

The New Metrics of Influence and Visibility

The more AI Overviews show in search and the more AI Assistants are adopted, the less relevant rankings and all the metrics associated with them become. In the world of LLMs it's more about whether AI systems use your content to form an opinion or answer. Which means we now need to be considering metrics like:

- **Share of AI Citation** How often your brand, product, or content is surfaced in AI generated answers across AI search including Google AI Overviews and AI Mode, ChatGPT, Copilot etc. in relation to other sources also being cited in the answers.
- **Paragraph-Level Visibility** In AI search sometimes small parts of your content might be extracted and cited independently. Tools like **seoClarity**, **Profound**, **Seer Interactive**, and experimental models from **iPullRank** are exploring this.

- **Entity and Topic Associations** Your brand or business may be mentioned (even without a link) in association with specific topics, features, or terms across trusted sources. As well as actual links or citations this is important to measure to understand the brand exposure you are getting and to assess what topics you are considered most relevant to.
- **Brand Citation in Multimodal Interfaces.** YouTube transcripts, TikTok captions, Reddit threads, and podcast metadata is scraped and synthesised by AI search models too and so visibility in these spaces is increasingly influential. Tracking your brand on these platforms is an important visibility metric and something that's been crucial for a while given how long they've been around!
- **Zero-Click Decision Impact.** This is really the next step up in terms of measurement. Understanding and reporting on whether your brand can influence purchase or consideration without a user visiting your site. This will be a particularly interesting one to measure once agentic shopping becomes more commonplace.
- **Conversation Inclusion** Finally, are your products, data, or brand narratives being referenced in real LLM-powered conversations, whether that's through direct queries or citation in outputs?

I've created, what I hope is a helpful, table considering existing performance metrics and new ones to try to give an idea of the changes that are likely to take place and why.

Metric	Type	Why It Matters Now	Still Valid?	Best Used For
Impressions	Traditional	Indicates exposure in SERPs or AI-Overview or increased demand	Yes	Useful in understanding improved visibility and demand
Clicks	Traditional	Measures active user engagement and intent to visit a site	Conditionally	Clicks are still relevant in various different instances e.g. in deep-research or product discovery queries
Sessions	Traditional	Measures active user engagement and intent to visit a site but also overall site visits and trends	Conditionally	Helps track broader traffic patterns from all sources - particularly referral from AI search!
CTR	Traditional	Shows effectiveness of a listing in traditional search and indicates changing patterns in SERPs	Yes (but context matters)	Still useful for traditional listings; less so in AI answers

Metric	Type	Why It Matters Now	Still Valid?	Best Used For
Average Position	Traditional	Tracks ranking position in linear SERPs at query level	Conditionally and Decreasing	AI interfaces break the notion of “position” or “rank” and personalised results make this more and more redundant
Engagement Rate / Bounce Rate	Traditional	Measures how well your content performs once visited	Yes - Conditionally	Still key on owned content, especially if linked via AI and when used with other metrics
Scroll Depth / Time on Page	Traditional	Gauges content engagement depth	Yes - Conditionally	Valuable for long-form or paragraph-targeted content and when used alongside other metrics
Share of AI Citation	AI-Specific	% of AI answers where your brand is cited or linked	New	Essential for tracking explicit visibility in AI Overviews, Perplexity, etc.
Conversation Inclusion	AI-Specific	Your brand is used in AI reasoning even if not cited	New	Shows latent visibility and brand salience in LLMs
Share of Voice in AI Interfaces	AI-Specific	Your % share vs competitors in AI-generated answers	New	Great for competitive benchmarking in generative UIs
Paragraph-Level Performance	AI-Specific	Tracks which paragraphs or sections get used in AI answers	New	Useful for structuring extractable, answer-ready content
Citation performance over time	AI-Specific	How often your brand reappears in AI responses over time	New	Indicates brand authority and trust persistence
Multichannel Visibility (TikTok, YouTube, Reddit)	Cross-Platform	These platforms are now source material for AI	Yes	Crucial for brand presence in AI-cited community content
Brand Mention Trends	Cross-Platform	Tracks the growth of mentions across media	New	Early signal for trending visibility and narrative growth
Sentiment-Weighted Mentions	Cross-Platform	Measures the tone (positive/negative) of brand mentions	New	Helps assess reputational health within AI training sources
Training Set Inclusion	AI-Specific	Whether your content is part of the LLM’s data corpus	Rarely visible	Important if known, but hard to measure directly

Beyond Clicks – Measuring Presence, Impact, and Influence

In my opinion, traditional SEO metrics still matter in specific circumstances — but the metrics that we report on need to adapt and change. In the age of AI search, your brand can appear in answers, guides, and product suggestions without a link and even without a brand mention - sometimes just your content alone may be used as part of an answer.

I can't stress enough how important unique measurement frameworks are now. The data shows that adoption of AI search differs across industry, that it is used more for specific types of search or conversation **and** that referral traffic from AI search differs in quantity in different niches.

Depending on the industry you are in, the types of content you are producing, the intent of the searches being made and many variables we don't know about individuals - clicks may still be an important metric in some cases, share of voice in another or frequency of brand mentions the next.

Define clear objectives and put the right measurement framework in place to track performance against them.

 **KPIs are evolving but the foundation of good measurement doesn't change.**

5.3 Tailoring Your Strategy to the Platform – Because Not All AI Search Is Built the Same

No AI search platform is the same - their mechanics are all slightly different, they provide different answers and they work better for specific tasks. They retrieve, interpret, and present content differently. The sources they cite most frequently, the training data they're exposed to, the way they structure answers, and the prominence they give to citations all vary.

That means the tactics that work well for one may be less effective on another - particularly when it comes to the sources that they pull from. James Allen wrote a brilliant article over at search Engine Land, 'How to get cited by AI: SEO insights from 8,000 AI citations' using data from [Rankscale.ai](https://www.rankscale.ai). The article and the data give a really detailed insight into the different sources that each AI search platform favours and specific websites and types of websites they cite the most.

Below I've included a very top level overview of the differences between ChatGPT, Perplexity and AI Overviews in terms of what they excel in, what they're not so good at and the sources they cite the most.

AI search Platform	Pros	Cons	Trusted Sources
ChatGPT search	Real-time searches, conversational, excellent for long-tail and very specific prompts	Struggles with local prompts and accuracy - especially with dates and travel planning - not a great experience for commercial queries	Favours trusted, authoritative and factual sources like Wikipedia, Reuters, Financial Times etc. Avoids UGC (User Generated Content) and forums
Perplexity	It's transparent and cites sources for every response, great for summarisation and information on news related topics	Frequent incorrect answers and hallucinations, image and local searches	Values expert and trusted sources and review sites - mainly cites blogs/ editorial content e.g. Investopedia. Some UGC sources
AI Overviews	Surfaces a broad range of media, very good at commercial and transactional searches	Poor on-going interaction, less good at very long-tail queries,	Pulls from the broadest range of sources including blogs, well known news sites, community content, forums (like Reddit), social media, favours deep pages over home pages.

The information in the table is the tip of the iceberg when it comes to the nuances of each AI search platform - particularly as it only includes three of them. But, it makes the point I want it to as we move into talking about key strategic initiatives for AI search. Just as audience, conversation and demand data are important, so is understanding the landscape you are optimising for.

Your brand's presence in AI search may not be evenly distributed — and that's OK. Optimise for where your audience is most likely to engage. Understand each platform's strengths, limitations, and citation behaviours — and tailor your strategic efforts accordingly.

Implication: One-size-fits-all doesn't work.

To build the right AI visibility, your strategy must diversify across formats, tone, and distribution. For example, getting cited in Perplexity may require insightful, expert-led content and strong presence in trusted specialist review sites. Whereas, to perform well in AI Overviews, if you've had a forward thinking approach to SEO over the last few years and evolved alongside it, you may already be performing quite well...

5.4 Brand Building & Digital PR – Fuel Awareness, Mentions & Authority in AI Search

Two studies, one done by [Ahrefs](#) and another by [Kevin Indig](#) demonstrate that the highest correlating factor with high AI search visibility is Branded Web Mentions and Branded search Volume. i.e. the number of times your brand is mentioned across the web and the demand for your brand.

This makes sense. Where traditional search engines have always used links as a measure of authority, LLMs draw from a vast corpus of web documents and build an understanding of your brand from words in that text. So, if your brand appears frequently – especially in association with positive sentiment and specific topics – it tracks that it's more likely to be mentioned in a favourable context and alongside those high-frequency themes because consistent, topical mentions reinforce topical relevance and authority.

→ How do brands increase web mentions?

Firstly, let's refer back to 'Deep Audience Research' and performance discovery because, without understanding where your audience hangs out and what they like, you're not going to build campaigns that resonate or encourage brand conversation. Once you have that - although links aren't 'exactly' the same - there is a huge correlation in the process that goes into acquiring or building citations and mentions - in this case, what we call 'Digital PR'.

Why It Matters

Digital PR's role in traditional SEO has always been about building authority and increasing brand touchpoints - because visibility extends far beyond your own domain, And, in the world of AI-powered search I think it becomes even more important. Mentions in trusted media, inclusion in roundups, appearances on podcasts, references in research papers, proprietary data sets and research papers... they all build brand authority and increase the likelihood of inclusion in AI-generated answers.

What To Do

→ Proprietary Research & Data Reports

Original research, data studies, or whitepapers are more valuable than ever. If your site is the source of data – survey results, market trends, specific pieces of research and stats – then you're offering information that not only the model, but the people using the model, can't easily find elsewhere. This makes your data uniquely citable - it also makes the likelihood of a click higher.

→ Earned Placements & Podcast Guest Spots

Put yourself or your people forward for inclusion in industry roundups, podcast interviews, or as the voice of expert commentary - they serve as really powerful trust signals and AI engines often interpret these as deeper validation. Especially if they are in the right sources for your target AI search platform.

→ 📰 **Newsjacking with Authority**

One of the fastest ways to gain authoritative mentions in high-authority media is to respond quickly to relevant news stories with relevant commentary, quotes, or data.

→ 📚 **Thought Leadership Series**

This is something that the SEO community are good at. There are several really good examples of excellent blogs that create multi-part content series that are structured into clever topical themes. These are often referenced externally (like in this whitepaper) and as a result are visible in more documents across the web, helping to build long-term brand narratives and authority.

→ 📊 **Indexes, Rankings & Reports**

Things like “Top 10 vegan cities in the UK” or “City-by-city AI adoption index” naturally attract media attention, links, and citations because it gives news outlets a headline and something to talk about. This differs from proprietary research and data reports in that these indexes can use pre-existing data sets to create a new story. But, if your brand is, or becomes, the owner of the dataset, that’s even more powerful for AI visibility.

→ 🗓️ **Cultural Moment Tie-ins & Reactive Campaigns**

These campaigns have to be quick moving and cleverly link to broader cultural, seasonal or calendar events — like Pride Month for example or, potentially, a viral TikTok trend - I heard Yetis are big right now?! These campaigns often earn engagement and press coverage that feels more organic and human - something LLMs are more likely to prioritise than generic brand content.

→ 📖 **Evergreen Resource Hubs**

Creating and maintaining high-quality resource pages, glossaries, or explainers (e.g. “The ultimate guide to EV charging” or “How is AI changing search?”) provides extractable, high-trust content that LLMs can reuse. (I guess the proof will be in the pudding).

→ 🛍️ **Product PR & Showcase Campaigns**

Getting your product featured in lifestyle roundups, recommendation lists by experts, or hands-on reviews are great for building strong credibility signals. These kinds of placements are good for conversions but they are also likely to be used as AI training data. If your product is repeatedly mentioned with positive sentiment, clear descriptions, and category relevance, it increases your chances of being cited or recommended in AI-generated answers (especially in prompts with purchase or comparison intent).

→ 🗣️ **UGC (User Generated Content) & Review-Rich Campaigns**

User-Generated Content — from TikTok unboxings and YouTube reviews to Instagram testimonials and Reddit discussions—increasingly feeds into the data pool that LLMs reference. This kind of content brings authenticity and relatability that traditional brand campaigns often can’t. Actively seeding and amplifying high-quality UGC, incentivising reviews, or collaborating with niche communities can help shape how your brand is perceived and surfaced in AI-generated answers.

For example:

- Verified product reviews on retailer sites
- Authentic mentions from influencers
- Reddit threads organically recommending your service
- TikTok or Instagram users showcasing real use of your product

This isn't just good for conversion — it's good for machine-readable trust. As we've established throughout this piece, AI systems build answers from loads of different online content — sometimes without linking back to or citing the original source. Instead, they recognise and surface brands that are consistently mentioned in publications they 'view' as trustworthy. That's why Digital PR is no longer just about backlinks - arguably it hasn't been *just* about backlinks for quite some time. It's about building brand authority and being part of the credible, high-quality conversations that AI assistants use to shape their responses.

Digital PR is a must-do. We often find it is one of the lesser requested services that is harder to sell-in to clients but it is so valuable and will become even more so. So, build unique assets, create campaigns that are worth talking about and earn credibility through mentions to position your brand to be surfaced in AI-powered search results.

Increasingly, the quantity of mentions and visibility in the right places is key and so is what those mentions say and the sentiment behind them. Visibility with negative sentiment can be harmful. This is obvious, right? And it has always been the case. Word of mouth advertising has long been one of the most positive forms - if you have a good product, people will talk about it favourably and others are more likely to use it. On the flip side, if people are saying negative things about your brand it should naturally be something that is addressed, never mind its impact on visibility in AI search. Yet so many brands don't do this - it's actually quite startling.

Not only is it important IRL (in real life) it's also key when AI assistants don't just index content but synthesise opinion, a brand's reputation isn't just a side note — it's a signal. One that LLMs, just like humans, interpret, remember, and use to inform how you're presented in future responses.

Which brings us neatly to the importance of Reputation Management.

5.5 Reputation Management – Shaping Brand Perception in the Age of AI

Reputation today is no longer shaped just by traditional media, direct customer interaction or formal reviews: consumers share their experiences across so many different platforms now, including Reddit, TikTok, YouTube, Glassdoor, podcasts and even formal and legal disputes. These are the spaces where narratives form and gain traction and these sources feed directly into the systems that inform AI-powered search.

Large language models don't form opinions – they synthesise responses and reflect opinions and sentiment that is already out there. If your brand is consistently talked about in a negative way – or misrepresented without you knowing – this can influence how you're cited [or even excluded] in AI responses.

So, managing user-generated content (UGC) no longer just sits under a brand or social purview. Reviews, discussions, and organic mentions are visibility signals in AI search. Managing them proactively – and encouraging the right kinds of contribution – has become critical to how your brand is understood and represented across the evolving search landscape. It's an AI search imperative.

Key Actions:

- **Track conversations about your brand across AI-visible platforms** (Reddit, Trustpilot, Glassdoor, TikTok, Instagram, YouTube comments etc.)
- **Respond proactively to negative press or community conversation.** NB - LLMs also scrape historical content so doing this retroactively is also important
- **Encourage and build positive, authoritative content** that outweighs poor-quality or outdated mentions
- **Use schema markup** (e.g. sameAs, Review, Person, Organization) to link content to trusted sources
- **Keep your Google Business Profile current** with events, special offers, up-to-date photos etc. - basically treat it as a social media channel.
- **Reviews and Q&As are key**, and require continual monitoring and engagement to build trust in your brand.
- **Encourage genuine positive reviews** on sites like Google, Facebook, and Yelp
- **Have a crisis plan ready** in case you need to execute comprehensive reputation repair campaigns across search, social, press releases, etc.

But visibility alone isn't enough.

Once you've tackled brand perception, the next critical layer is how *trusted* and *verifiable* your own content and presence are to AI systems. It's not just about who's mentioning you – it's about whether those mentions meet the credibility standards LLMs rely on. That's where Trust Signals come in.

5.6 Trust Signals – Earning Your Place in the AI Ecosystem

As LLMs and AI become leading elements of the search experience, trust, authority, credibility, and security become even more fundamental. Increasingly, being trusted will be about being considered credible enough that language models will cite, after being competitively evaluated, across a wide range of other content. We've spoken about external perception sentiments through Digital PR and Reputation Management – but what about how brands present **factual**, **trustworthy**, and **secure content** about themselves in owned spaces.

Why It Matters

- Just as in traditional search, Google's AI Overviews demote (and even exclude) content that fails to meet E-E-A-T (Experience, Expertise, Authoritativeness, Trustworthiness) standards
- Studies suggest that [some LLMs actively avoid citing controversial or low-quality sources](#), particularly in more sensitive YMYL (Your Money or Your Life) niches
- Traditional search trust signals including author bios, publication dates, references, backlinks (citations) etc. now also contribute to 'AI credibility layers'

What To do

- **Build a reputation as a reliable source.** Publish content that's reviewed, sourced, and fact-checked.
- **Ensure that the content and claims you publish are accurate** and align with standards, **especially in regulated industries such as health, finance, and legal.**
- **Monitor trust signals on key pages**, including author credentials and expert quotes (more on this in the next section).
- **Create in-depth content** that provides unique information not found elsewhere
- **Cite reputable sources** and include outbound links to authoritative websites - this is especially important when publishing content that is within medical, financial, legal, and technical spheres. Where relevant, display references or bibliography as this helps both users and machines verify accuracy.
- **Use clear, concise language** that's easy for both users and LLMs to understand
- **Use conversational language** that includes synonyms and related phrases that reflect how people naturally speak and search.
- **Create comprehensive FAQs** that directly answer common questions, and address multiple related 'follow-up' questions.
- **Format your content** with lists, tables, and bullet points to enhance its appeal.
- **Ensure your content is secure** by serving it via HTTPS with up-to-date certification, publish clear privacy statements, and user-friendly T&Cs.
- **Demonstrate content freshness** by adding review timestamps, clearly marked with "Updated on [date]" so that both users and models can recognise freshness.
- **Create author profiles and biographies** to illustrate the experience behind the articles being written.
- **Promote expertise by using schema markup** for articles, authors, organizations, reviews, etc., to help crawlers (and LLMs) understand and trust your content.

Demonstrating trust and credibility is essential — but if AI systems can't easily parse or verify your content, those signals may go unseen.


That's why it's important to ensure your content isn't just trustworthy to humans, but also legible and deemed credible by machines — through structured data, semantic clarity, and precise formatting that enables extraction and citation at scale.


5.7 Semantic Optimisation and Structured Data - Optimising structure for machines


Machine readability is fundamental as we move into a world of search underpinned by LLMs. It's not just about satisfying crawlers or — unfortunately — putting users at the heart of what you write — it's about ensuring that your content is *comprehensible*, *extractable*, and *citable* by these new systems.

AI systems rely heavily on content that's clearly structured, demonstrates clear relationships between entities, and includes good semantic markup in order to better build responses. This is the layer that connects **E-E-A-T** signals, a brand's **reputation**, and **content value** to the actual mechanisms of AI visibility.

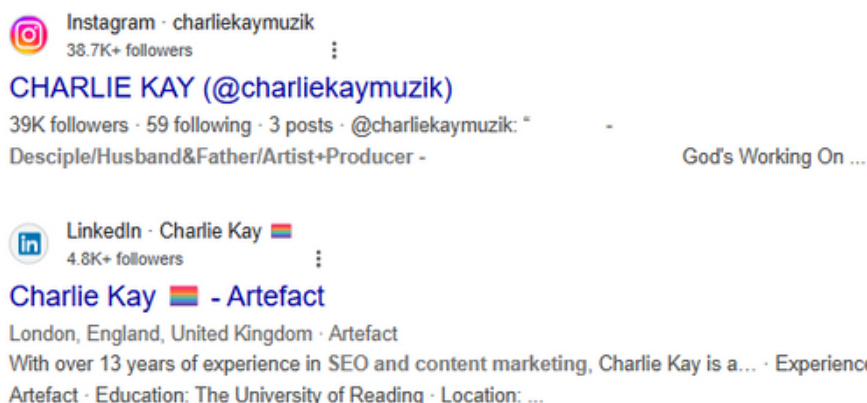
Why It Matters

 **LLMs don't 'crawl' pages the way traditional search engines do.** Instead, they take in content, tokenise it, encrypt or 'vectorise' it, and learn the relationships between different words, sentences and concepts using different systems. This often involves pulling from structured signals and patterns to understand what's being said, who it's coming from, and what it means in context.

 **AI Overviews and assistants rely on content that is structured well** in order to extract facts, attribute credibility, and also to determine relationships between entities - including concepts, brands, and people.

 **Semantic structure** helps models better distinguish between entities, for example "Charlie Kay the Head of SEO" vs "Charlie Kay the musician". Proper use of **<Person>**, **SameAs**, or **Author** schema with LinkedIn/Wikidata links helps resolve ambiguity about which person the article is about. Another example would be "Apple" the fruit vs "Apple" the company. Using contextual schema (**Product**, **Organization**, or **Food**) and internal linking (to other tech-related content vs. nutrition guides) will help models to resolve the intended entity.

[The image below (at the top of the next page) shows two listings for 'Charlie Kay' in Google search Results. One is for 'Charlie Kay' the musician and the other is for 'Charlie Kay' Head of SEO at Artefact. Using structured data can help to make the two entities more distinguishable to machines.]



🔗 **Structured data improves visibility in both traditional and AI-powered search**, supporting traditional rich SERP features including carousels, featured snippets, product displays, and now – increasingly – LLM-generated answers. In fact, recently at SMX – search Marketing Expo in Munich – Fabrice Canel confirmed that **schema markup helps Microsoft's LLMs understand content**.

What To Do

✅ Implement Schema Markup

Structured Data or Schema markup is additional code that can be added to your page in different formats (for example Microdata, or JSON-LD which is Google's preferred formatting) to provide additional rich information about your content – you can read more about **structured data markup in Google search here**.

Here are some examples of commonly-used Schema markup, with your schema being unique to the content that you have, the type of business you run, etc.

- **Article, NewsArticle, BlogPosting** (for editorial content)
- **Product, Offer, Review** (for commerce)
- **FAQ, HowTo, QAPage** (for instructional content)
- **Person, Organization, SameAs, Author** (to promote identity and trust)
- **Event, Recipe, Video, Breadcrumb** (more examples depending on content type)

🔧 If you or your development team want to know more, tools like **Schema.org**, **Google's Rich Results Test**, **Dentsu's Schema Generator** and **JSON-LD Playground** can help to generate valid markup.

✅ Manage Visibility with Preview Controls

Google enables site owners to control what appears in both traditional listings and AI-generated experiences. You can use preview controls like **nosnippet**, **data-nosnippet**, **max-snippet**, and **noindex** to better manage how your content is displayed. **N.B. Setting more restrictive permissions will limit how your content is featured in AI Overviews and other AI surfaces.**

✨ Make Entities Clear - How Search Understands Real-World Meaning

As discussed, AI systems use entities to better understand 'real-world' meaning. By clarifying these entities in your content you can improve the chances of accurate retrieval, attribution, and citation in AI-driven search results. You can do this by:

- ✓ Using internal links and consistent naming conventions to reinforce relationships
- ✓ Adding sameAs schema to show clear connections between digital identities and profiles
- ✓ Referencing external sources (e.g. Wikipedia, LinkedIn) to validate key people etc

Examples where this matters:

- Similar product names: e.g. iPhone 14 vs. iPhone 14 Pro Max
- Overlapping brands: e.g. Crucial (as a brand) vs. Micron (as parent company)
- Authors with the same name: e.g. Charlie Kay the SEO vs. Charlie Kay the musician
- Writers with content across domains: Ensuring AI recognises the same author on Medium, LinkedIn, and your blog

👁 Extractable Formatting & Paragraph-Level Answers

AI Overviews and LLMs extract specific paragraphs, not whole pages, which is a complete departure from traditional search that ranks a list of links to... you've guessed it, pages.

To ensure better extractability in order to increase the likelihood of your paragraphs being used within answers, you should structure content to:

- Summarise key points at the top of the document
- Ensure each section has a clear subheading (<h2>, <h3>, etc.)
- Keep paragraphs short, standalone, and well-labelled
- Give paragraphs a clear objective and purpose
- Answer questions clearly in 2–5 sentences, and in bulleted lists where possible

💬 This will increase the chance that your content is *extracted as a response*, not just featured as a reference.

🧩 Use Semantic HTML

More than ever, the actual structure of content matters. The use of proper semantic tags for example <article>, <section>, <aside>, and <main> helps to make page layouts meaningful to bots, by breaking down content into named sections.

This helps LLMs (and assistive tech) interpret page layout, intent, and – crucially – the importance of different elements. A lot of websites often overlook or misuse semantic tags, and rely on generic <div> or tags which offer no real semantic clarity.

🍷 Ensure Clean, Crawlable Code (more on this in crawlability)

For content to be shown, cited or linked to by AI search, their bots have to be able to find and digest it. While some LLMs generate responses purely from pre-trained data, many of the systems powering AI search today — like Google's AI Overviews, Perplexity, Bing's Copilot, and Gemini — actively crawl and index the live web to retrieve real-time content. This means that clean, crawlable code still matters for visibility.

A little teaser below, with more on this in '[Crawlability and AI Agent Access](#)'

- Avoiding Javascript-heavy or SPA (Single Page Application) designs that obscure content
- Ensuring key content isn't hidden behind logins or interstitials
- Using canonical tags properly
- Making sure robots.txt and meta tags don't block important pages

← END In this new world, structure is a source of meaning for LLMs and AI search. Schema markup, semantic HTML, and well-optimised page architecture all help LLMs understand what content is about, who it's from, and when and why it matters. Content **must** still consider the end user, but now it's also about making your brand's knowledge machine-readable, unambiguous, and citable.

'Semantic optimisation' is the link between content, credibility, relevance, and AI visibility. If it's done correctly it increases the likelihood of your information being surfaced and reused in AI-powered search.


But the web isn't just made of words. As AI systems become increasingly multi-modal and capable of interpreting images, video, audio, and alternative platforms, visibility depends on more than just semantics and structure. Brands now need to think beyond the page — ensuring that every content format, and every platform it appears on, is accessible, discoverable, and understood by both users and machines.

🖼️ 5.8 Multi-Modal and Multi-Platform Optimisation – Make All Your Content Accessible

As I hope I showed when deconstructing its evolution, search hasn't been a 'text-only' experience for quite some time. Pinterest is image-first and YouTube is video-first, and both these media formats have been embedded in traditional search since ~2007!

Today, the models that AI search is built on are multimodal by default. They extract meaning from **images, videos, tables, diagrams**, and more. So, if your content isn't optimised or structured across these formats, then it may not be discovered or won't be adding as much value as it could.


Why It Matters

 AI systems provide answers in a range of different formats and surface content that helps answer a user's query in the clearest and cleanest way possible. This is also true of traditional search, where structured formats like tables and charts make content more extractable and often more trusted – think optimisation best practices for featured snippets.

With help, models like GPT-4o and Gemini can now actually 'see' and understand images. If you want your content to be recognised and used in answers, then you need to consider the way that images and video are embedded on the page, using things like descriptive alt text, adjacent captions, and structured data.

For example, but by no means limited to, content wrapped in semantic tags like ``, `<picture>`, `<figure>`, `<figcaption>`, or properly structured `<video>` blocks with contextually-linked text is far more likely to be understood and surfaced in AI-generated responses than media buried in divs or injected via JavaScript.

- `` is used to embed an image onto a web page
- `<picture>` is used to provide multiple sources for an image, allowing a browser to choose the best one based on device capabilities
- `<video>` is used to embed a video onto a web page
- `<figure>` is used for self-contained content that is relevant to the main flow of the document, but can also be removed without affecting the document's meaning
- The `<figcaption>` element is used to provide a caption or title for the content within a `<figure>` element

 Well-labelled visuals support answer synthesis and paragraph assembly, especially when paired with explanatory copy and context.

What To Do

Honestly, if SEO has been a part of your marketing strategy, then all of the following are things that you should already be doing and if you aren't then it's time to get your foundational elements in order:

✓ Use Descriptive Alt Text and Captions

Include contextual descriptions that explain what the user is seeing and why it matters. For images, describe the visual content and pull out any embedded text. For videos, include a short summary or caption that outlines the key points, topics covered, and any spoken dialogue that is not otherwise represented in the surrounding text. This helps AI models understand what the video contains and why it's useful - this is also important for accessibility.

✓ Image & Video Sitemaps

Submit dedicated image and video sitemaps (or include `<image:image>` and `<video:video>` in your main sitemap). These act as direct signals to search engines and AI agents, guiding them to hidden or dynamically loaded media. In sitemaps you're also able to include additional metadata in sitemaps i.e. location, caption and title etc. to provide more information about the images.

✓ Serve Media via Crawlable, Accessible HTML

Where possible, don't paste an image of a chart – use structured formats and semantic HTML like `<table>`, `<figure>`, `<picture>`, and `<video>` or even CSV downloads where applicable, so AI systems can extract and interpret the data semantically. For videos that display charts or stats, consider providing a transcript or adjacent text summarising the data shown.

✓ Structured Data

We've touched on structured data previously but it would be remiss not to mention it again here. Using structured data for images and video helps search engines of all types understand more about your media.

✓ Label Visual Content Clearly

For infographics, charts, and videos, use headers, legends, or callouts to clarify what's being shown. Use descriptive file names (e.g. `laptop-comparison-2025.mp4`) and surrounding text that aligns with your key topics or product themes. This increases the chance your content is matched to relevant prompts and supports high cosign similarity between your media and specific topics.

✓ Content-Linked Metadata

Tie visual assets to a topical or entity-based structured data – for instance, connecting a product image with `Product` schema or an event promo video with `Event` schema. This increases the chances of multimodal comprehension and reuse.

✓ Transcript Markup for Video & Audio

Include clean, structured transcripts for all audio and video content, either inline or via structured markup (e.g. using `Text`, `Transcript`, or `MediaObject` schemas). This supports indexing, retrieval, and summarisation in AI systems.

✓ Add Key Moments or Video Segmentation

Just like YouTube's "Key Moments", break longer videos into chapters or segments using 'Clip markup' (part of `VideoObject`). This helps LLMs and AI search systems extract the most relevant portion for specific queries.

Example: A travel video could include time-stamped segments like "Where to stay", "Local transport", "Food tips".

✔ Use Thumbnail Optimisation Strategically

Thumbnails serve as visual anchors for both human users and AI-powered interfaces. Use unique, content-relevant thumbnails that visually summarise a piece of content/media's purpose. Use the `thumbnailUrl` field in your VideoObject, Image schema etc. It influences selection in AI-generated carousels and will potentially be used as a citation image for specific content.

✔ Use <track> for Subtitles and Captions

For video elements, use the `<track>` tag with `kind="captions"` or `kind="subtitles"`. This not only improves accessibility but gives AI models another entry point to the video's semantic meaning – particularly in multilingual contexts.

✔ Use MediaReview Markup for Explainers or Fact-Checks

If your video debunks or reviews media (e.g., reaction videos, explainer content), consider using the MediaReview schema. This is an emerging area of interest for AI and fact-checking ecosystems.

🔍 Platform Discovery Optimisation – Show Up Where Audiences Are searching

AI search ingests and surfaces content from a broad ecosystem of platforms across the web – including social channels, video platforms, forums, and thought leadership hubs. In fact, starting [10th July, Instagram is going to open up its content to be indexed in Google](#).

To earn visibility, brands need to show up in these ecosystems.

🔧 Practical Ways to Engage by Platform - a few examples:


Reddit → Join or monitor relevant subreddits - if you're doing this as a brand I'd recommend responding to objective questions e.g. 'Does [brand] sell [x] product?' rather than answering more subjective queries about the brand where answers will look biased e.g. 'Are [brands] [products] the best ones?' Answer FAQs, run expert AMAs, and build trust via consistent presence. Reddit is a high-weight source for Gemini and Perplexity.

YouTube & TikTok → Create video content that gives users an insight into your brand and answers questions they are asking straight from the source. Far too many brands use these platforms to repost adverts - this is fine as long as it's mixed in with genuinely helpful, authoritative and expert content. You can optimise specifically for these platforms by creating answer-friendly videos that are clearly spoken, well-titled how-tos, explainers, product demos etc. Use captions, timestamps, and descriptive summaries.

LinkedIn, Medium, Substack → These platforms are regularly used as sources in AI Answers, particularly for B2B queries so being active in these spaces is key. If it's applicable, brands should develop a regular posting strategy on these platforms, publishing thought leadership pieces, company and product updates, opinion pieces etc. with clear authorship and citations.

Quora, and other niche forums → Many LLMs pull from alternate community sites or forums to Reddit where specific, user-generated Q&A or visuals help clarify niche topics. There are millions of different forums on the web and many within specific niches. Search out relevant forums in your niche, understand whether they are cited in AI search results and learn how they work. Being in these spaces will help you better understand your audiences and the conversations and pain points they have, it will help with encouraging positive sentiment and visibility in AI search.

Business Listings and Merchant Centre → Ensure your local, business and product content is up-to-date, especially for commerce, travel, and retail brands.

 **Takeaway:** *Visibility in AI search doesn't start with your website – it starts with where users share, search, and shape the narrative. If you're absent from those ecosystems, AI might overlook you entirely.*

5.9 UX & Content Architecture – Designing for AI Discoverability at Scale

While structured data and semantic markup help machines parse individual pages and paragraphs, the broader architecture of a site – how content is organised, linked, and experienced – plays a critical role in how LLMs interpret authority, relevance, and usefulness across an entire domain.

You can think of this as 'experience engineering' because it is for both **humans and machines**. When users and AI bots land on your site, they should be able to quickly understand who you are, what you offer, and where to go next to easily find what they want. For LLMs trained on patterns and associations, this clarity increases the chances your content gets cited, extracted, and reasoned with.

Why It Matters

Clear content architecture improves context and relationships.

LLMs look for topical consistency and entity associations, and structuring your site well will reinforce your expertise. This has been relevant for quite some time, as traditional search behaviours have changed resulting in much longer form queries. It has made sense to move away from keyword level optimisation to think more topically and ensure a site demonstrates expertise in areas that are relevant.

📦 **Modular content architecture supports paragraph-level citation.** Breaking up information into digestible, semantically consistent paragraphs makes it easier for AI systems to pull out relevant sections for responses.

🔗 **Intentional internal linking structures help LLMs trace topical authority.** Content that's interlinked with relevant subtopics helps reinforce domain-level expertise and increases visibility.

📱 **Mobile-first and accessible layouts matter.** Up to 59.45% of all web traffic globally comes from mobile devices. At its core, mobile-first design prioritises easily navigable pages, a smooth, fast and flexible experience, progressive enhancements and the prioritisation of essential content - all of which are key for AI search optimisation.

⚡ **Fast, Responsive Pages Improve Access and Visibility.** Speed still matters. AI systems often crawl or preview pages in real time. So if content times out or lags, it might be skipped. Ensuring that all key pages load quickly, especially on mobile is vital.

👉 **A fast-loading page improves both crawl success and user feedback signals, which can reinforce inclusion in AI-generated responses.**

What To Do

✅ Use Scalable Page Templates

Design repeatable formats for articles, guides, product pages, and FAQs that use consistent heading structures, content modules, and CTAs (Calls To Action).

✅ Structure Content Hubs & Pillars

Group related content around themes highly relevant to your brand or business to show depth. This is particularly helpful for LLMs when forming topical associations and reasoning chains. If you have a category landing page about a specific topic, then you could create a sub-folder level editorial hub – including carefully designed internal linking – to help build topical authority in this area.

✅ Design with 'Extractability' in Mind (see section above)

Make sure each section can stand alone:

- Use descriptive subheadings
- Summarise key points at the top of each section
- Keep paragraphs focused and succinct

✅ Simplify Navigation

Avoid deep nesting and overly complex menus. LLMs and users both prefer flatter, more intuitive site structures. Use breadcrumb trails, cross-linking, and contextually relevant CTAs.

✓ Optimise Mobile Experience

Use responsive design with a properly configured viewport, ensure fast load times, legible fonts, intuitive navigation, mobile-friendly spacing, prioritise important content in page structure, appropriately size touch targets, and ensure compatibility with all device orientations and screen sizes.

✓ Optimise Page Load

This is an highly complex and important topic and to do it justice would need a entirely new whitepaper. I've outlined some of the key elements below:

- **Core Web Vitals:** Helpful as a simple prioritisation framework for load performance.
- **Code & Asset Efficiency:** Minimise, defer, and streamline code, scripts, and assets.
- **Image & Media Optimisation:** Compress, size appropriately, and lazy load where needed.
- **Server & Delivery Speed:** Optimise server response, use CDNs, and efficient caching.
- **Render Prioritisation:** Prioritise above-the-fold content and visible elements.

🕸 5.10 Crawlability & AI Agent Access – Don't Block the Bots That Matter

Many AI systems — from search-integrated LLMs like Google's AI Mode and Bing's Copilot to standalone tools like Perplexity and ChatGPT — now **actively crawl the live web** to retrieve and synthesise real-time information. Patrick Stox published a helpful post June 18 2025 called 'Meet the New Web Crawlers: AI Bots Are Closing in on search Engine Bots' - worth a read if you want to understand more.

This means your site's visibility in AI search isn't just about what's indexed by Googlebot — it's also reliant on whether these **AI-specific bots** can access your content as well.

But many brands and businesses are knowingly and unknowingly **blocking** the very agents that power next-gen discovery experiences.

Why It Matters

🤖 **LLM-based tools like** ChatGPT (via Browse), Perplexity, Claude, Google's AI Overviews and AI Mode **rely on access to web content** via dedicated bots (e.g. GPTBot, PerplexityBot, ClaudeBot, etc.)

🔒 **If your site blocks these bots** — whether via robots.txt, firewalls, or rate-limiting — **your content might never appear in AI-generated answers**

🔍 **Indexability and crawlability issues can lead to poor visibility** in both traditional and AI search, even for otherwise high-quality content

🧠 **You can't influence AI answers if your content isn't accessible in the first place**

What To Do

✓ Audit bot access regularly

Ensure key bots are not blocked by [robots.txt](#), Cloudflare/firewall settings, or server configurations. Bots to check for:

- GPTBot (OpenAI / ChatGPT Browse)
- Googlebot / Google-Extended
- ClaudeBot / Anthropic-User
- Bingbot / CCBot / Copilot Agent
- PerplexityBot / Perplexity-User

✓ Whitelist AI Bot IPs in Firewalls

Bots like GPTBot may be blocked at the server or WAF (e.g. Cloudflare) level. Consult OpenAI's crawler IP list and Anthropic's documentation to allow by IP range.

✓ Render Essential Content Server-Side

LLMs and most bots still struggle with JavaScript-heavy content. Use server-side rendering (SSR) or pre-rendering to expose meaningful content. Avoid hiding key assets behind JS, tabs, interactivity, or client-only rendering frameworks (e.g. React, Vue).

✓ Avoid Blocking Valuable Content with Meta Robots

Make sure high-value content (resources, product pages, help articles, blog posts) is available to bots, even if it's gated for regular users.

✓ Avoid JavaScript-only rendering

AI systems may not render content that relies entirely on JavaScript — especially SPAs (Single Page Applications). Use pre-rendering or SSR (Server-Side Rendering) to ensure content loads.

✓ Use canonical tags smartly

If duplicate or similar content exists, use canonical tags to point AI systems to the preferred version, reducing dilution and confusion during synthesis.

✓ Avoid “noindex” and “nosnippet” on valuable pages

These tags may prevent pages from being surfaced or cited in AI responses. Don't use noindex on key content pages. Don't apply nosnippet if you want content to appear in AI answers — it removes paragraph-level visibility for AI Overviews and AI Mode.

✓ Optimise internal linking for discoverability

Strong internal linking ensures crawlers find your deeper pages, and reinforces content relationships for semantic understanding.


💡 **AI visibility starts with accessibility. If your content isn't crawlable, it isn't findable — and if it isn't findable, it won't be surfaced or cited in AI search.**


5.11 Feedback Loops & Model Monitoring - Track, Test, Learn Across AI Ecosystems


You can't optimise for what you don't understand — and AI systems don't give you 'keyword rankings' or even prompt level data - this isn't even split out by Google in search Console, the data is there but aggregated with that from traditional search. To understand performance and stay visible, we need a new kind of feedback loop. That means regularly testing how your brand appears across different AI platforms, identifying gaps or misrepresentations, and feeding those insights back into your strategy.

There are already lots of tools that monitor performance in AI search, from referral traffic to prompt level visibility, to citation share of voice and mention frequency. Nikki Lam at search Engine Land put together a helpful table [comparing the different available AI search tools](#).

Why It Matters

 As we've discussed at length, AI search is different to traditional search and so is monitoring its performance - the metrics and KPIs change.

 AI search is dynamic and evolving — meaning your presence (or absence) may shift weekly, based on model updates, content changes, or user interactions. It's important to keep on top of this.

 As well as using new tools to track performance in this new landscape, I'd recommend running manual proactive tests, logging outputs, and segmenting results by query type, platform, and citation pattern.

What To Do

✓ Build a Monitoring Matrix

Track presence across ChatGPT (Browse + GPT-4o), Gemini, Perplexity, Claude, Copilot - or whichever platforms the tool you are using provides/ are most important to you because that's where your audience is. Monitor answers for brand mentions, product recommendations, and citation patterns.

✓ Log & Analyse Prompt Outputs Over Time


Log prompts, responses, and citations over time. Segment results by topic, query intent, and funnel stage (e.g. discovery, comparison, conversion).

✓ Set Up Brand Visibility Benchmarks

Define visibility goals for each platform. Is your brand being mentioned? Are you being cited in explanations? Included in recommendations? If other brands get more mentions, why?

✓ Feed Learnings Back into Strategy

If AI systems are citing competitors, investigate why — and use those insights to improve your content structure, digital PR, or trust signals.


 This isn't one-and-done work. AI-generated results are probabilistic and prone to drift. Your job is to keep tuning and testing.


5.12 Agentic Interaction Readiness - Optimise for Decision-Making (Without Clicks)


As AI Agents shift from search assistants to action agents — booking travel, choosing products, resolving support requests — brands must rethink what visibility means. Even for bottom of the funnel queries, you're not just trying to earn a click anymore. You're trying to earn a place in the agent's decision-making model.


- *Sidenote - A couple of consideration points.*
 - a. 'Google Shopping' was launched in 2012 and 'free listings' (released to help small businesses during the pandemic) were released in 2020. So, shopping in search has been evolving and becoming more complex ever since - meaning some of what is required for Agentic Checkout isn't new.
 - b. It'll take some time to get it totally right and before agents are allowed to actually make purchases for users at scale. There are a lot of governance risks and legal implications to consider.

Why It Matters: Some Examples

 An agent might choose a product based on specs, availability, and trust — not product copy.

 In travel, it might recommend hotels or itineraries based on structured data and reviews - this already happens somewhat.

 In customer service, it might auto-answer a user query using your Help page — or a competitor's.

 These interactions happen without a user ever landing on your site — but the outcome (purchase, booking, brand choice) still hinges on whether your data was available, structured, and trusted.

What To Do

✓ **Optimise Product and Service Content for AI Agents** Ensure product specs, availability, price, reviews, and shipping info are structured and up to date. Use schema like Product, Offer, and AggregateRating.

✓ **Prepare Support and FAQ Content for Agent Extraction** Use FAQPage, HowTo, and clear HTML structure.

✓ **Cover Full Use Cases, Not Just Keywords** Agents don't match queries — they interpret intent. So optimise for scenarios like "best baby monitor for travel" or "affordable hotel near King's Cross with parking".

✓ **Ensure Interoperability** Use open data formats and APIs where appropriate — especially in industries where agent-led bookings are emerging (e.g. travel, retail).

Final Takeaways

AI Search is Already Here

AI has been part of search for a long time, and brands that have been following this evolution will (or should) already have been making iterative shifts to their strategies. From keywords to content clusters and from a focus on the brand website to a multi-platform strategy.

Now AI search is here. Google's AI Mode, a fully conversational AI search interface, is rolling out as part the world's largest search engine — first in the US, with other markets close behind.

A Defining Shift

AI search is growing rapidly and is [projected to overtake traditional search by 2028](#). AI has been reshaping how people search, discover, and make decisions for years and now it will even more so. Critically, AI search now changes how answers are constructed — moving away from simple indexing and ranking towards reasoning, synthesis, and conversational outputs. This shift is already transforming the nature of visibility in search.

For Right Now: Traditional Search Still Dominates

Despite the pace of AI adoption, traditional search remains dominant. Google still holds the majority of global search share, and most search traffic and visibility continues to come from the classic search experience. AI search adoption is growing quickly, but right now, it still makes up only a small fraction of total market share and is concentrated within specific industries, intents, and demographics.

And, AI platforms — including Google's AI Overviews, AI Mode, Bing's AI answers, and ChatGPT browsing results — still rely heavily on traditional search indexes. Strong performance in classic SEO closely correlates with citations and inclusion in AI-powered search (despite what some might say). Visibility in one space still reinforces visibility in the other. **Modern SEO isn't dead — it's simply evolving.**

Looking Forward

Different Platforms, Different Priorities

It's essential to remember that AI search platforms don't all work in the same way. They draw on different types of training data, they prefer different answer formats, and they retrieve from different sources with varying frequency. Optimising across them demands platform-specific awareness — this is not a one-size-fits-all approach.

The Power of Data and Measurement

Perhaps the most important thing you can do now is to use your data wisely. It's essential to deeply understand where your audiences are, whether they're using AI search, which platforms they trust, and what kinds of content resonate with them across those spaces.



Your measurement frameworks needs to evolve to fit the new landscape and provide the insights needed to help you move forward.

Modern measurement should be highly granular, capable of flexing across platforms, and pivot to track emerging AI-metrics like citations, passage-level visibility, and multi-modal appearances. Not all content should be measured the same way. Some content will still drive clicks. Some will surface in AI-generated answers without ever mentioning your brand. Some will influence decisions without generating a single session on your site. Capturing the right data will help you better demonstrate success and feed it back into a more successful, ever-evolving strategy.

A Final, Final Word

The brands that will win in this new ecosystem are those that stay visible across platforms, structure their content for both human and machine readability, evolve their measurement frameworks, and consistently prioritise real value, trust, and usability.

The future of search is layered, multimodal, and agentic — and the time to get ready is now.



Strategy

&Transformation

- Λ Data & AI Strategy
- ΛData & AI organization
- ΛData Maturity Assessment
- ΛCorporate Training
- ΛHackathons
- ΛData & AI Days
- ΛGenAI Academy
- ΛArtefact AI Summits



AI

Acceleration

- Λ AI & Gen AI Factory
- ΛData & AI for Operations
- ΛAI for Customer Care
- ΛData & AI for Private Equity



Data

Foundations & BI

- Λ Data Governance & Management
- ΛData, New BI and Self Business Intelligence
- ΛData for Sustainability



IT

&Data Platforms

- Λ Data-centric IT
- ΛData Platform
- ΛCustomer Data Platforms
- ΛCloud Services and Certifications



Marketing

Data & Digital

- Λ Consumer Data Environment
- ΛMeasurements (MROI) & Insights
- ΛData Valorization & Category Management
- ΛData-driven Sales
- ΛMarketing Analytics
- ΛCX & Digital Marketing
- ΛGMP Certified Reseller
- ΛMedia Services & Certifications



ARTEFACT

AI IS ABOUT PEOPLE

ARTEFACT UK

1st Floor, 62 Queen Street
London EC4R 1EB

CONTACT

hello - uk@ar tefac t.com
+44 (0) 207 802 9500

artefact.com

