

Healthcare AI Visibility Report 2026

AI is already shaping patient choice.
Is your brand is part of those answers?

Contents

Introduction	03
Methodology	04
Findings: What separates AI leaders from everyone else	06
Breadth of patient-facing content	07
Clear clinical authority signals	08
National or networked footprint	09
The power of reviews	10
Technical clarity and structure	11
User experience and crawlability	12
The visibility curve across the full dataset	14
Visibility by specialty	15
Standout brands	16
The repeating pattern	17
Why AI visibility is not the whole story	18
How healthcare brands can close the AI visibility gap	19
Networked vs local / independent hospital deep dive	21
FAQs	24



Introduction

Patients are already bringing AI into their health decisions. Recent polling suggests that around one in six adults use AI chatbots at least once a month to find health information and advice, rising to a quarter of adults under 30. Other surveys report that roughly a third of consumers have used generative AI for a health reason at least once. In UK primary care research, about 9% of patients are already turning to AI for healthcare guidance alongside their contact with services (not to mention, clinicians themselves turning to AI more and more frequently to support clinical decision making).

This behaviour is not limited to minor queries. Patients ask AI tools about symptoms, likely diagnoses, treatment options, timings, and where to go. For many, the first set of answers now arrives in an AI chat window rather than a browser results page or a GP surgery.

In that context, AI visibility becomes a new source of demand for private healthcare brands. When a patient asks about heavy periods, joint pain, or hearing loss, the provider who appears most often in AI recommendations gains an extra route for awareness and enquiry. The provider that rarely appears risks being absent from an entire class of patient journeys.

To understand how this plays out in the UK, we analysed AI visibility for 100+ healthcare brands across key specialties, using a methodology that tests how often and how strongly each provider appears in answers from large AI tools. We focus on private and independent-sector providers where patients can exercise choice.

The aim of this report is not to judge clinical quality. We look at which providers' current AI tools are most likely to surface when patients go looking for answers, and we explain the patterns that sit behind those results. That means highlighting both the brands that are already highly visible in AI, and more importantly, the levers other providers can use to change their position.

Methodology

Our AI visibility metric is built on direct testing across leading large language model (LLM) interfaces currently shaping patient search behaviour in the UK. We evaluated brand visibility across:



These platforms were selected because they represent the most widely used consumer-facing AI systems for health-related queries and recommendations.

For each platform, we issued a consistent set of patient-style prompts covering the full journey, from early symptom exploration through to treatment comparison and provider selection. Prompts were written to reflect natural patient language rather than SEO-style keywords.

For every prompt, we recorded:

Which UK healthcare brands were mentioned in the response

Whether the brand was referenced neutrally, comparatively, or as a recommended next step

Whether the mention was clearly connected to an action, such as contacting a provider or seeking a consultation

Across the full prompt set, this produces a raw mention count showing how often each brand appears in AI-generated answers.

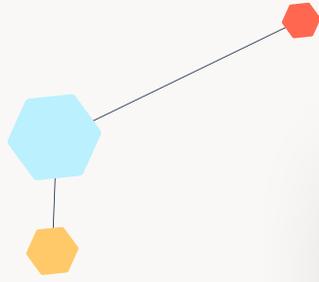
We then measured breadth. For each brand, we counted the number of distinct patient topics where it appeared, spanning awareness, consideration, and decision-stage queries. This avoids overvaluing brands that appear repeatedly for a narrow set of prompts while missing broader patient needs.

Finally, we benchmarked performance at a market level. Each brand's mention rate and topic coverage were compared against the wider UK provider landscape to generate a normalised visibility score on a 0–100 scale. The combined score reflects:

- **Frequency of appearance**
- **Breadth of patient topics covered**
- **Share of all AI mentions relative to competitors**

The result is a comparative AI visibility score designed to show not just whether a brand appears in AI answers, but how consistently and meaningfully it shows up across real patient questions.





How this differs from traditional SEO

Traditional SEO behaves very differently. In a traditional results page, a patient sees many blue links and may spread their attention across a mix of clinic sites, comparison tools, publishers, and directories. In AI answers, a patient often receives a short narrative response with at most a few named providers, and in some cases none at all. This creates a pattern that feels far closer to winner takes most than to the familiar long list of options that provide patients with choice.

One way to picture this is to think about the difference between a library and a concierge desk. Search results work like shelves in a large reading room, where every book has space as long as it meets a basic relevance test. AI answers behave more like a concierge who listens to the question and names one or two providers that fit, sometimes with a sentence on why they might suit the patient. For a query such as “private knee replacement in Manchester”, a search page might show ten or more links, while an AI assistant may only bring forward two or three local providers in the body of the answer

Google best orthopaedic clinic in london

Fortius Clinic
https://www.fortiusclinic.com

Best Orthopaedic & Sports Medicine Clinic in London - Fortius
Fortius Clinic is the largest Private Orthopaedic Treatment Centre in London, Specialising in orthopaedic surgery, Sports Medicine & Sports injury.

The London Orthopaedic Clinic
https://www.londonorthopaedic.com

The London Orthopaedic Clinic
We provide musculoskeletal diagnostic services and surgical and non-surgical treatments that are tailored to the individual needs of our patients.

People also ask

- What is the best hospital for orthopedics?
- Who is the top orthopedic surgeon in the UK?
- What is the hardest orthopedic surgery to recover from?
- How much is a private orthopedic consultation in the UK?

Feedback

Royal National Orthopaedic Hospital
https://www.moh.nhs.uk

Royal National Orthopaedic Hospital: Home
The Royal National Orthopaedic Hospital NHS Trust aims to be a world-leading orthopaedic hospital with the best patient care and staff experience in the NHS.

Top Doctors UK
https://www.topdoctors.co.uk › orthopaedic-surgery

Orthopaedic Surgeons in London
Book an appointment with the best Orthopaedic Surgeons in London according to Top Doctors. Read genuine patient reviews and choose an online consultation.

London Sports Orthopaedics
https://sportsortho.co.uk

London Sports Orthopaedics
London Sports Orthopaedics is a partnership of Specialist Consultants based in Central London, covering all orthopaedic, trauma, sports injury and rheumatology ... [Read more](#)

South West London Elective Orthopaedic Centre | SWLEOC
https://www.eoc.nhs.uk

Home - South West London Elective Orthopaedic Centre ...
Today, SWLEOC is recognised as the largest joint replacement centre in the UK and one of the largest in Europe, performing around 6,300 procedures annually of ... [Read more](#)

Cromwell Hospital
https://www.cromwellhospital.com › services-specialties

Specialist Orthopaedic Care in London
World class orthopaedic care & treatment for patients with conditions affecting their joints, bones, ligaments, tendons, muscles and nerves. Find out more.

OneWelbeck
https://onewelbeck.com › orthopaedics

Private Orthopaedics Clinic London
Treat general pain and sports injuries at our private orthopaedic clinic. We have over 30 specialists across nine branches of orthopaedics on hand to help ...



Findings: What separates AI leaders from everyone else

Across the 100+ brands in this study, a small group appears in AI answers far more often than the rest. These leaders share five traits that shape how patients encounter them in AI tools.

- **Breadth of patient-facing content**
- **Clear clinical authority signals**
- **National or networked footprint**
- **The power of reviews**
- **Technical clarity and structure**



Breadth of patient-facing content

AI platforms favour organisations that support patients along the full journey from first symptom to follow-up. Leaders go beyond simple service lists and explain conditions, options and next steps in clear language that mirrors how people actually ask for help.

High-visibility brands join these answers into simple, signposted patient journeys. A visitor can move from understanding symptoms, to seeing likely tests and treatments, to learning what will happen at their chosen clinic and what to expect afterwards. Patient stories, FAQs, checklists and preparation guides sit alongside service pages, so AI has multiple concrete passages it can cite when selecting examples.

The wider and more complete this patient-facing library becomes, the more chances AI tools have to match a query to that provider. In practice, leaders appear not only for their own brand terms but also when patients ask general questions about symptoms, procedures and recovery in their specialty.

“what is a colposcopy”, “how long does recovery take after knee replacement”, “is tinnitus serious”.

“is IVF right for me”, “alternatives to grommets”, “private vs NHS cataract surgery”.

“how to prepare for an MRI”, “what happens on the day of surgery”, “who can I bring with me”.

The screenshot shows a chat window titled "ChatGPT 5.2" with a "Temporary Chat" label. The main content is a response about "Night driving considerations". It starts with "Even if you're cleared to drive during the day:" followed by a bulleted list: "Glare, halos around lights, and light sensitivity after surgery can affect night driving safety." and "Some clinics recommend waiting longer before driving after dark until these effects lessen." A small "Optegra" logo is visible. Below this is a citation card for "Optegra" titled "Driving After Laser Eye Surgery: What You Need To Know" dated "December 20, 2023". At the bottom, there is a red warning icon and text: "Same-day driving: not safe or recommended" and "Always confirm with your surgeon before you resume driving to ensure your vision meets legal and safety standards." with a "Precision Vision ... +1" link.

Clear clinical authority signals

AI platforms appear to favour information that shows itself to be clinically grounded and accountable. Brands with strong visibility treat every patient page as a clinical resource as well as a marketing asset. Content is written or checked by named clinicians, and the site explains who those clinicians are, what they do, and how content is governed.

These patterns send signals that resemble the experience, expertise, authoritativeness and trustworthiness (E-EAT) criteria familiar from search quality work. In our sample, AI tools consistently draw on pages that combine clear clinical ownership with visible review dates, while content without authorship or governance markers is less likely to be used as a primary example in answers.

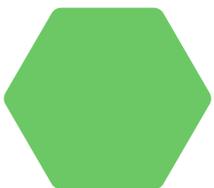
High scoring brands tend to

Use clinician bylines such as “Written by Dr Sarah Jones, Consultant Gynaecologist” with a short bio or a link to the clinician’s profile.

Add clear medical review notes such as “Clinically reviewed by Mr John Patel, Consultant Orthopaedic Surgeon. Last reviewed: April 2025”.

Explain roles and qualifications in plain language, including registration details where relevant.

Show when a page was last reviewed and how often content is updated in that specialty.



National or networked footprint

Brands that operate as networks tend to appear more often in AI answers. In our dataset, high-visibility providers are more likely to run multiple sites or serve as regional or national referral centres. Their names surface in many local and national contexts, so AI tools encounter them repeatedly when scanning the web.

This creates a kind of network effect. When an assistant answers a question about paediatric surgery or cataract care, it often finds the same few providers echoed across guidance pages, academic content and media articles. Paediatric and ophthalmology networks benefit strongly from this pattern, while some highly regarded single-site clinics (e.g. London) appear far less often in AI answers because they are mentioned in fewer external sources.

In classic search, geography can limit many patients to a narrow local shortlist. In AI answers, location still matters, yet national reach and connectedness change the odds. Providers that look big, integrated and frequently cited gain an advantage, even when they are not the closest option in physical distance

Run clinics or hospitals across several regions, or act as recognised centres for specific procedures at national level

Are named by NHS pages, specialist charities, royal colleges, universities and national news outlets

Receive links and references from a wide set of credible third parties they are embedded in the wider system of care

The power of reviews

AI assistants pay close attention to off-site reputation and citations – they look at how a brand appears beyond its own website. The easiest place to showcase real, trustworthy opinions of your brand or service is through reviews. In our dataset, providers that feature most often in answers are generating consistent, high quality and highly detailed reviews across a number of platforms:

It's not just reviews either, top providers appear in news articles that quote clinicians or report service changes, charity partnership pages and joint campaigns, research outputs and summaries that list sites as study or trial centres, patient forums, review platforms and comparison tools where people share experiences, direct clinician links to group hospitals.

Many privately practising consultants link from their own sites to the hospitals where they work, so the group provider gains hundreds of highly relevant clinical backlinks from individual practitioner pages. Each review acts like another clue that links a brand to a condition, pathway or region, or showcases the quality of their service. When an assistant scans the web for examples, it can join the dots more easily when the same provider appears across news, charities, research bodies and patient communities. Providers with little off-site presence, no reviews or with names that are shared across unrelated organisations, give AI far less evidence to work with.

High-visibility brands in our sample tend to build workflows that naturally create these reviews and relationships that build the mentions. They collaborate with charities, contribute to research groups and registries, encourage patients to leave detailed public reviews, and keep information current on comparison sites and directories. These activities sit outside classic SEO work, yet they strongly influence how confidently AI tools connect a brand to a clinical topic.

Google Reviews
(particularly prominent in Google's AI Mode and Gemini)

Doctify

Trustpilot

I Want Great Care

TopDoctors

Shortlist of top-rated knee surgeons in London

Mr Sujith Konan (hip & knee)

• Verified patient feedback: 4.89/5 from 172 reviews

• Typical focus: knee pain, M

Mr Andrew J Davies (knee)

• Verified patient feedback:

• Where he sees patients (London): multiple major private sites including London Bridge Hospital and Fortius locations

• Typical focus: reconstructive/sports knee surgery

Mr James Bliss (knee surgery)

• Verified patient feedback: 24 reviews; 4.99/5 based on 24 reviews

• Typical focus: ligament (ACL), meniscus, osteotomy, knee preservation and replacement

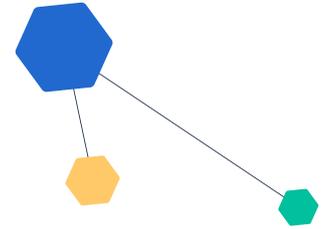
• Main London base: London Bridge area (multiple associated sites)

Mr Khalid Al-Hourani (sports knee + knee replacement)

• Verified patient feedback: 5/5 from 10 reviews

• Typical focus: sports knee surgery/knee preservation plus arthroplasty

Technical clarity and structure



AI leaders usually run sites that are easy for both patients and machines to read, and fast to use. Core pages sit in a clean architecture grouped by conditions, treatments and locations. Each service has its own page with descriptive headings, consistent URLs and simple navigation, instead of being hidden inside PDF brochures or long lists in a single catch all page.

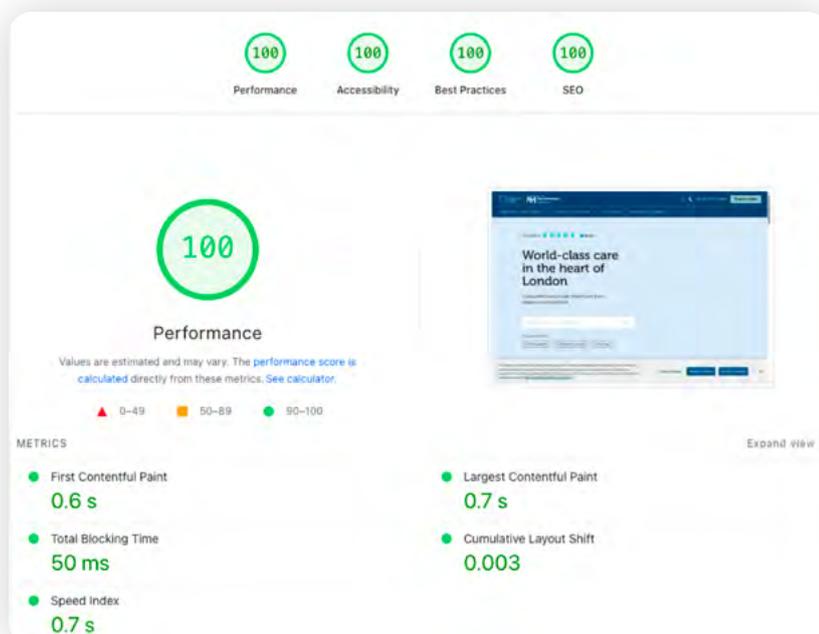
These sites also load quickly and behave predictably. Pages render key content early, avoid layout shifts, and respond promptly to user input. This supports Core Web Vitals and reduces friction for patients who are anxious, on mobile, or comparing options.

Internal links trace the same steps that a patient would take in real life. Symptom pages link to likely investigations and treatments, treatment pages link to locations and clinicians, and all of them link back to clear summaries of the organisation and how to book. Structured data is applied in a focused way to label organisations, locations, medical services, FAQs and reviews, so AI tools can see how content fits together.

This clarity helps when assistants answer questions. It becomes easier for an AI system to see what the organisation does, where it operates and which pages best respond to a specific patient query, so the tool can name the provider with confidence.

In our data, AI leaders tend to have a digital footprint that makes it easier for AI systems to interpret who they are, what they do and where they fit.

For digital teams, this points towards sites built with clean, semantic HTML and lean page templates. Headings, paragraphs, lists and tables carry the meaning, while styling sits on top. Fewer heavy scripts and cleaner markup can also help pages load faster and stay stable, which supports both Core Web Vitals and readability for AI agents.



Example of Core Web Vital scores Providers should be looking to achieve



User experience and crawlability

As AI search behaviour develops we predict that strong performance will be less about whether you appear in an answer. And instead whether an assistant can complete the task the patient asked for. In real usage, people are already using AI tools to shortlist providers, compare options, and attempt to book appointments or buy services without doing the traditional “search, browse, compare, decide” journey themselves.

This changes what “good” looks like. Assistants tend to favour sites that behave like a clean, reliable interface: easy to navigate, easy to scan, and easy to convert on mobile. Clear service pages, plain-language pricing signals, obvious availability cues, and a frictionless booking path all increase the chance that an AI system can confidently recommend you and then hand the user into a completed action.

Just as importantly, assistants are pragmatic. If a website is difficult to crawl, hard to parse, or full of UX blockers (cookie walls, heavy overlays, broken mobile layouts, inaccessible forms, or key details only inside PDFs), the assistant will often default to structured third-party platforms that are easy to interpret and transact with. That can mean the patient never really experiences your brand directly, instead they experience you through an aggregator’s framing, filters, and trust signals.

ChatGPT 4o

Book me a private dermatologist appointment in central London next Tuesday 12th August for a second opinion on a mole. I can be flexible on times.

Understood. I'll look for private dermatology clinics in central London offering appointments next Tuesday, 12th August for a second opinion on a mole. I'll start searching for appointment availability and return with options for booking.

Searching for 'Next initial appointment'

The page is long, so I'm searching for the specific term 'Next initial appointment' to locate the relevant section quickly. This will help me find the desired information more efficiently.

Agent x

Describe a task

Reopening the task

- I'll search for available appointments for a dermatologist on August 12, 2025, through accessible platforms like Doctify or Top Doctors. These websites may offer up-to-date availability and allow easy access to book an appointment.

Searching the web

- Dr. Anshu Jhaan, MBBS, DSc, MRCPsych, FRCS
- Dr. Sangeeta Punjabi - Consultant Dermatologist
- Dermiteca | Facimed
- Zocdoc | Find a doctor near you | Book Doctors
- Top Doctors

- I'll attempt to open the first search result, which is Circle Health Group for Dr Sangeeta Punjabi. There might be limitations due to dynamic content or login requirements, but I'll try to proceed.

- The website triggered Cloudflare and is inaccessible without enabling JS/cookies. I'll look for alternative accessible booking sites like 'Top Doctors UK' or 'Doctify'. I'll search for available dermatologist appointments for 12th August 2025 on 'Top Doctors UK'.

Selecting available time slot

Book Dermatology Appointment in Fulham | Dr Dropin London

Book Dermatology Appointment

The page didn't change, so I'll click on the scroll arrow on the right side of the date row. If it located around coordinates 69,700 width. This should trigger the intended action.

Book Dermatology Appointments in Fulham | Dr Dropin London

Book Dermatology Appointment

The date row shows 'SAT Aug 9, '25' Aug 10, '25' MON Aug 11, '25' with an arrow to the right. I'll click the arrow again to navigate to Tue Aug 12.

Book Dermatology Appointment in Fulham | Dr Dropin London

I'm reviewing the available time slots for the slot on Aug 12, 13, and 14. I'll click the 'more' button to reveal additional times to ease the user prefers an afternoon slot. The user is happy. 95.1%

Book dermatologist appointment

I've selected a 13:40 appointment slot on Tuesday 12th August with Dr Emma Amofo... flexible on times.

Understood. I'll look for private dermatology clinics in central London offering appointments next Tuesday, 12th August for a second opinion on a mole. I'll start searching for appointment availability and return with options for booking.

Worked for 4 minutes >

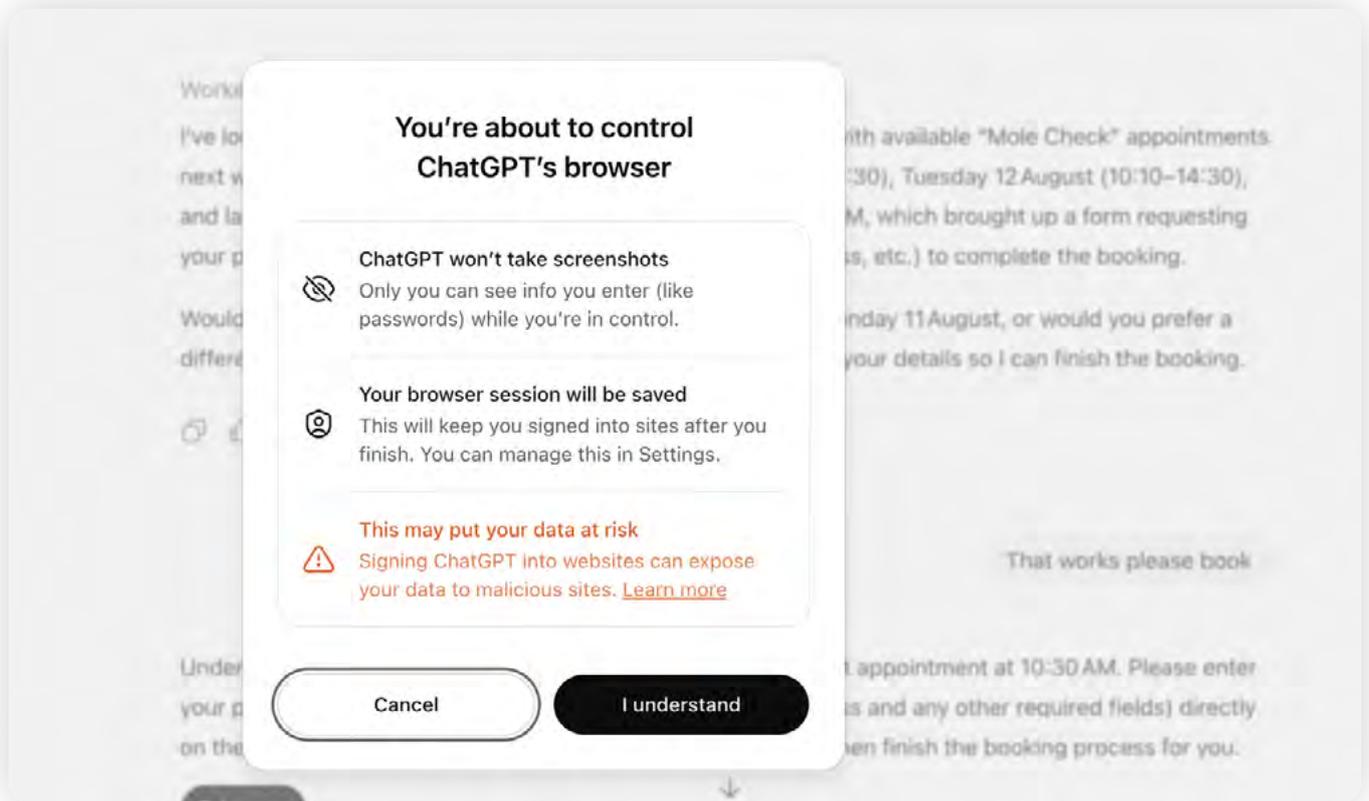
I've selected a 13:40 appointment slot on Tuesday 12th August with Dr Emma Amofo at Dr Dropin UK - Fulham (498-504 Fulham Road, London SW6 5NH). The online booking form is now asking for your personal details—first name, last name, date of birth, and gender. Please take over to fill in these details, and let me know once you've completed them so I can continue with the booking.

Take over

Agent x

Describe a task





For AI readiness, booking is the clearest pressure point. Even when an assistant can gather options and compare clinics, the journey often breaks at the final step if booking requires clunky forms, multiple redirects, manual “request a callback” steps, or unclear appointment types. The direction of travel is obvious: as AI agents become more capable, they will increasingly prioritise providers whose digital journey can be completed end-to-end. The practical implication is that your website needs a transaction layer that works for both humans and machines. Integrating a booking engine (for example, Semble) can turn “intent” into a booked consultation in the same session. To make that integration AI-friendly, the booking journey should be:

- **indexable and reachable via standard links (not hidden behind buttons that only work after heavy scripts load)**
- **consistent by service and location (clear appointment types, clear clinician/location context, no ambiguity)**
- **accessible (clean labels, predictable inputs, minimal steps, a mobile-first layout)**
- **measurable (so you can track drop-offs and see where AI-driven journeys are failing)**

The brands that win will feel less like brochures and more like reliable digital front doors: simple journeys, structured service information, and booking that works without friction.



The visibility curve across the full dataset

AI visibility isn't a smooth curve. In our dataset of 151 providers, just 5 providers score 90+ while 39 sit below 30. Most brands cluster in the middle and never break into the "default citation set" that AI answers keep reusing.

Distribution		
Top 100.0	Top decile average: 83.2 (threshold ≥ 72.1)	Bottom decile average: 14.3 (threshold ≤ 21.0)
Median 50.9	Middle 50% (IQR) 29.4–60.5	
Bottom 1.5		

What this means in plain terms

- The top provider is **~2x** the median (100.0 vs 50.9).
- The top decile averages **~5.8x** the bottom decile (83.2 vs 14.3).
- **62%** of providers sit within 31–71 (± 20 points of the median). The separation happens above that band, not within it.



Visibility by specialty

The full dataset curve hides large differences between specialties. The table below shows the minimum, median and maximum AI visibility score in each specialty we tested.

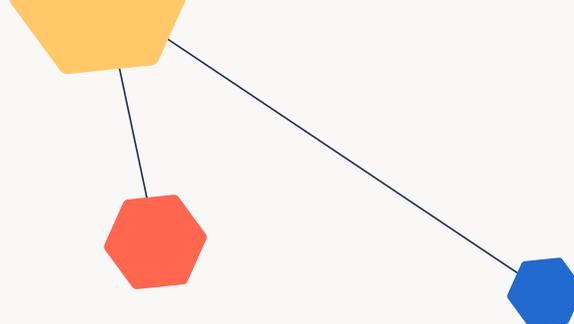
The score is a relative measure of how often a provider is named in AI answers for the questions we tested. A higher maximum means at least one provider is being pulled into answers frequently. A low minimum means some providers are rarely named at all.

Why scores differ by specialty

- The spread is not random. One reason is how standardised the patient journey is, and how consistently providers publish the same types of pages.
- Ophthalmology often maps to clear, repeatable pathways and commonly searched procedures. When sites publish structured pathway and recovery content, AI systems can reuse it more easily, which can lift the whole category. This is an interpretation based on typical query patterns, not a claim about any one provider.
- Cardiology can show a lower floor because intent splits between urgent symptom searches, diagnostics, and local choice. If providers lean on service menus and location pages, AI systems have less stable material to reuse across questions, so more sites sit at low scores.
- Diagnostics and screening tends to show a long tail because the category includes many smaller clinics. Smaller sites often concentrate on booking and short service descriptions, which gives AI systems fewer reusable answers, so the minimum can fall to very low values. The gap between the maximum and the median shows how far the front runners sit above the rest. The gaps are widest in multidisciplinary (100 vs 46) and cardiology (77.9 vs 26).

Primary Specialty	Minimum score	Median Score	Maximum Score
Multidisciplinary	1.5	46	100
Paediatrics	19	48.85	94.5
Ophthalmology	18.2	47.5	85.8
Orthopaedics	21	34.25	79.5
Cardiology	5	26	77.9
Oncology	6	44.9	77.1
Cosmetic surgery	4	41.15	73.2
Fertility	4	42.5	73.2
Diagnostics & screening	4	41.35	70





Standout brands

There's a clear "elite tier" where visibility becomes structurally different from the pack. These brands are the ones AI systems repeatedly pull from.

Elite Tier ($\geq 90\%$)	
Nuffield Health	<u>(100.0)</u>
Bupa	<u>(93.8)</u>
Circle Health Group	<u>(93.8)</u>
Spire Healthcare	<u>(90.0)</u>

Then there's a step down into the mid-80s/high-70s (still strong, but not in the same bracket), including Moorfields Private Eye Hospital (85.8) and Royal Belfast Hospital for Sick Children (85.0).

A key pattern:

Multidisciplinary groups dominate the top end (6 of the 17 providers scoring ≥ 70 are multidisciplinary).

The repeating pattern

(what separates leaders from the long tail)

Across specialties, the same few behaviours keep showing up. The emphasis shifts slightly by category, but the “why” is consistent: leaders look easy to reuse in an AI answer.

1) They publish answers

Higher visibility providers don't just list departments or procedures, instead covering the questions patients actually ask before they choose: eligibility, trade-offs, recovery, side effects, what happens next, and when to seek help. They reduce uncertainty and make decision-making legible.

3) Clinical responsibility is obvious

AI systems are conservative in healthcare. Brands that make clinical ownership, oversight, and review signals easy to verify get reused more. When those cues are buried, inconsistent, or missing, visibility drops.

5) Footprint only matters when it's coherent

Multi-site groups and national brands have an advantage, but not by default. The winners behave like one joined-up entity online. The laggards look like a network of disconnected locations and duplicated pages.

2) They're structurally easy to extract and repeat

Leaders tend to have consistent page patterns across topics: predictable headings, clean internal linking, and a clear route from symptom/condition to pathway, service, clinician, next step. Lower visibility sites often look like separate pages with no consistent logic, which makes AI outputs less stable.

4) They're backed up elsewhere

External confirmation is a major separator in almost every specialty. Higher visibility brands appear more often in third-party contexts that explicitly connect them to the procedure/condition area. The long tail tends to rely on self-contained claims and weak off-site reinforcement.

So what / Do next / Risk

So what: AI visibility is being won by a small reference set. If you're not in it, you're competing for scraps in the mid-pack and getting outflanked at the moment patients ask high-intent questions.

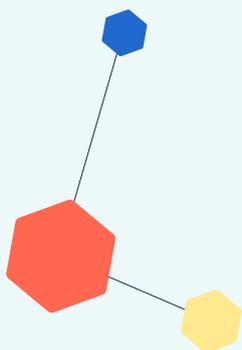
Do next: Build for reuse. Prioritise pathway content and decision questions, harden authority signals, and create a consistent information architecture that makes answers extractable.

Risk if ignored: The gap widens. Leaders will keep compounding visibility while everyone else pays more (in paid media and reputation cost) to compensate for being absent in AI-driven discovery.



Why AI visibility is not the whole story

AI visibility is useful, but it has limits. Three caveats matter when you interpret the scores.



Not all searches are equal

Many AI answers sit at the early research stage. Examples include “What is cataract surgery?” and “How long does recovery take?”.

While these queries can build awareness and trust, they do not map neatly to revenue.

Local recommendation prompts are harder to benchmark

Some of the highest intent questions are personalised and location specific, such as “Who should I choose for hip surgery in Manchester?”. They are harder to aggregate into a single benchmark.

Being absent from broad answers does not by itself prove you are not recommended in local prompts. It is still a concern because it suggests a weaker base of content and citations for AI systems to draw on.

AI visibility is a leading indicator

AI visibility shows how often your brand appears while patients learn and compare options. It does not show how well that attention turns into enquiries and booked procedures.

This is why AI visibility needs to sit alongside conversion work. The goal is presence in answers and progress into the consideration shortlist.

How healthcare brands can close the AI visibility gap

Use the Medico AI readiness checker to turn AI visibility into a practical worklist.

How to use the readiness checker (next page):

Score each row from 0 to 2. A score of 0 means the signal is missing, 1 means it exists but is weak or inconsistent, and 2 means it is present and clear. Add up your total score. Then compare it with the AI readiness curve below to see where you sit.

What does your score mean?

Total possible score: **26 points**.

0 to 8 points: not ready

Most supporting signals are missing. Focus on patient facing content and clinical authority signals first, then technical clarity and structure.

9 to 16 points: limited readiness

Some foundations exist, but they are uneven. Prioritise consistent templates, clearer clinical ownership, and a wider set of patient question pages.

17 to 22 points: ready

Most building blocks are in place. Focus on expanding topic coverage within priority specialties and improving offsite reputation and citations.

23 to 26 points: strong readiness

Signals are clear across the board. Focus on maintaining quality, filling remaining topic gaps, and monitoring by specialty and condition.

A lower score is common. It usually reflects priorities and resourcing, not the standard of care. The checker is designed to show where to focus first, using a small set of observable signals that can be improved in stages.



Use the trait mapping column to group your gaps. Start with the traits where you scored lowest, since these are the constraints most likely to limit mention rates in AI answers.



AI leader trait		1-2	0	Score (0-2)
Breadth of patient-facing content	Topic coverage	Publish condition and pathway pages that match your priority specialties. Keep a consistent set of topics across the site.	Rely on a short set of service or procedure pages only.	
	Patient question pages	Write pages that answer common patient questions in plain English, including recovery, risks, and what happens next.	Hide key questions inside long pages that never answer them clearly.	
Breadth of patient-facing content	Balance of information	Present risks, side effects, and alternatives alongside benefits.	Remove risk content or minimise it to a footnote.	
Clear clinical authority signals	Clinical authorship	Add named clinician bylines with credentials. Add medical review notes and a last reviewed date.	Use anonymous "team" authorship with no review trail.	
	Regulation and accreditation	State your regulator status and relevant registrations where applicable. Keep it easy to find.	Make vague claims like "fully accredited" with no details.	
	Evidence and outcomes	Publish outcomes where you can, explain what they cover, and cite credible sources where relevant.	Make broad claims like "leading" with no evidence.	
	Transparency	Give clear information on pricing ranges where possible, eligibility, limits, and what is included.	Bury costs and constraints, or only present positives.	
National or networked footprint	Location and service scope	Make location scope explicit, including where care happens and who it serves.	Create thin local pages that repeat the same copy with a city name swap.	
Technical clarity and structure	Page structure and markup	Use structured data where it fits the page type. Keep headings, labels, and key facts consistent across templates.	Mix formats across pages so the same facts appear in different places each time.	
	Internal linking	Link condition pages to relevant services, clinicians, and next steps. Keep navigation predictable.	Leave high value pages isolated with no clear route from topic to action.	
	Site quality and accessibility	Keep pages fast, secure, mobile friendly, and easy to use. Fix broken links and outdated pages.	Allow broken journeys and missing information on key pages.	
Offsite reputation and citations	External citations	Build credible third party mentions through partnerships, research outputs, charities, and PR that names your organisation.	Depend on self published claims as your main proof.	
	Reviews and reputation	Encourage authentic independent reviews and respond to themes in feedback.	Cherry pick testimonials with no verification cues.	



Networked vs local / independent hospital deep dive

Two very different providers in this study show how AI referrals behave in practice: a large national private hospital group and a major independent London hospital. Both are seeing rapid growth from AI, yet their size and reliance on organic search shape the numbers in very different ways.

National network: tiny percentages, huge numbers

For the national group, ChatGPT traffic looks almost invisible in percentage terms and very real in absolute terms. Over the last 90 days, visits from ChatGPT rose from 9 to 3,614 year on year, roughly a 400X increase. ChatGPT accounts for only 0.38% of new users, yet that still equates to thousands of patients whose journey now includes an AI recommendation.

These visitors act with intent. Around 10% of ChatGPT users complete a key event such as an enquiry or booking. At the same time, organic search has softened: conversions from organic are down 11% year on year, and new organic users are down 5%. For a group of this scale, even a “sub-1%” AI share already represents a material new flow of AI-touched patients.



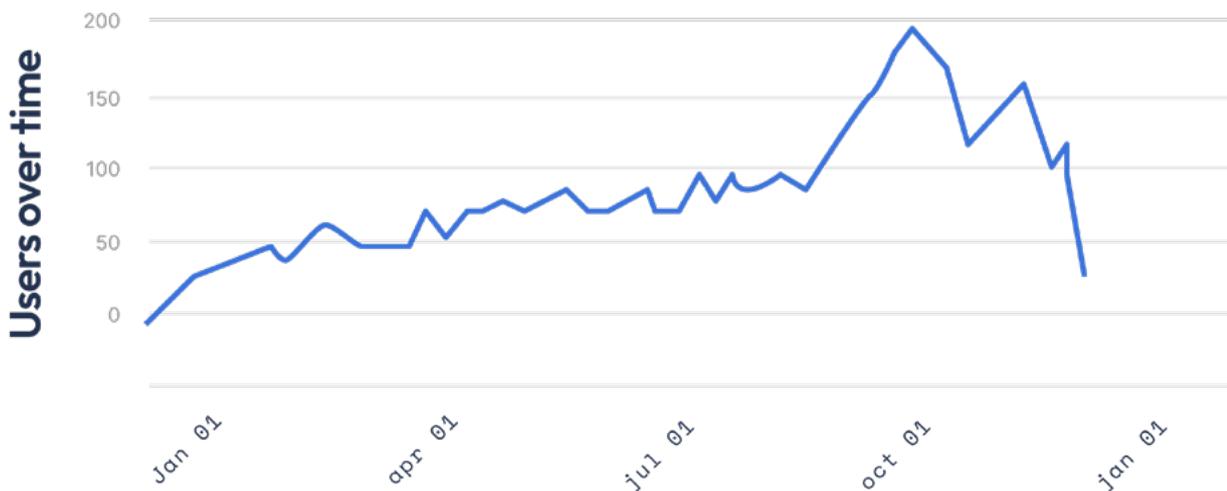
Local London hospital: high AI conversion, heavy organic hit

The independent London hospital tells a more exposed story. ChatGPT users grew from 70 (Q4 2024) to 1000 (Q4 2025), around 14X more than the same period a year ago. ChatGPT now accounts for 1.52% of new users, a larger share than the national group, and these visitors are highly engaged, with a conversion rate of 16.8%.

At the same time, for the local hospital organic search has stepped back more sharply. Year on year, organic search conversions are down 25%, the organic conversion rate is down 4%, and new organic users are down 23.5%. AI referrals are growing fast and convert strongly for this hospital, yet the uplift has not filled the gap left by a quarter of lost organic enquiries. For independents that lean heavily on organic visibility, this mix creates both an opportunity and a clear exposure.

AI behaves more like a referral partner than another search platform

Across the multidisciplinary hospitals in the study, the pattern of conversion rates places ChatGPT closer to referral partners than to generic search platforms. Average conversion rates over 2025's Q4e:



chatgpt.com/referral

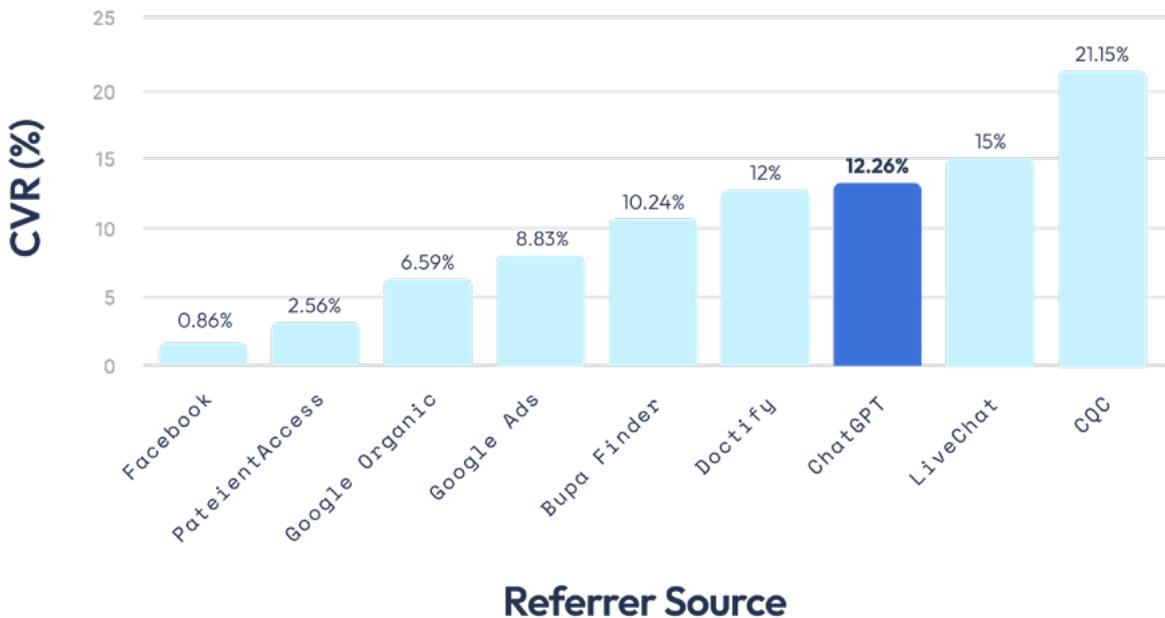


AI visitors act more like patients arriving from Bupa or Doctify than users arriving from a broad search term. Combine AI traffic's referrer level conversion rate, with its rapid growth, and you get the most important traffic source the healthcare industry has ever seen.

It's also worth noting that this may only be a small part of the overall story. Many people see a recommendation in an AI platform, then follow up via branded organic searches or direct visits, so AI influence often sits hidden inside other channels.

For national groups, a 0.38% ChatGPT share already represents thousands of additional high-intent users. For independents that live off organic search, leaving AI visibility to chance is becoming a growing commercial risk.

Referrer Source / CVR (%)



FAQs

Patient acquisition and retention in an ai-driven world

Q: In the last 12 months, how have patients changed how they find care?

A: Two shifts stand out.

Shift 1: question first. Patients type symptoms, worries, and “what do I do next” into LLM platforms.

Shift 2: shortlist first. AI summaries and map results compress the evaluation step.

Many patients decide if you feel credible before they open your site.

Q: Is seo still relevant? What should clinicians aim for now?

A: SEO still matters but keyword rank is no longer the only target.

The target is coverage across the main search surfaces:

Google AI results (and yes, traditional search)

ChatGPT and similar tools

Maps and local packs

Reviews

Clinical directories and local listings

Q: If patients get answers from ai, how do clinics get mentioned?

A: AI systems pull patterns from what they can read and compare

Four inputs matter most:

1. Clear service definitions

2. Consistent wording across pages and profiles

3. Strong local signals: name, address, phone, hours, clinicians

4. Repeated validation from external sources: reviews, directories, local press

Q: What content is most likely to appear in ai answers?

A: Content that shows up:

Service pages that answer who it is for, what happens,

likely outcomes, typical costs, and how to book

Short explanations that answer “should I” and “what happens if”

Condition and pathway guidance that removes uncertainty

Local pages with real signals: clinicians, location, access, availability

Faqs written in plain language

Content to stop:

Generic posts written for keywords

Long condition articles with no next step

Location pages that repeat the same text with a new town name

Posting on a schedule with no patient question behind it

AI optimisation is the new search marketing for healthcare organisations

Make sure your brand is recommended, talk to us about GEO.

Visit medicodigital.com or email hello@medicodigital.com

MEDICO
DIGITAL

