Winning Patients in AI Search

Idea In Short

Treat artificial intelligence [AI] search as a strategy problem, not a new channel or a vendor project. Work through four questions in order: can search systems read your content, why would an answer trust you, does the rest of the web agree with you, and how will you measure results. Fix eligibility first, rebuild your highest-revenue pages with named clinician input, align listings and profiles, and measure at the front desk and phone system, not through third-party scripts that may expose patient data under the Health Insurance Portability and Accountability Act [HIPAA]. Avoid shortcuts and volume-driven content. A practice that systems can read, trust and describe accurately is the one a generated answer names.

By Steve Bunker, CEO, Target Patients MD

Search is shifting from a list of links to a synthesized answer, and that shift changes what a medical practice is competing for. The goal is no longer only to rank; it is to be the source an AI-generated answer draws on and the practice it names. Practices that treat this as a new channel, chase shortcuts or measure the wrong things will waste a year. Those that treat it as a strategy problem, built on four questions about eligibility, authority, consistency and measurement, can compete on terms that favor the specialist over the biggest advertiser. Healthcare adds two constraints most industries do not face: a higher bar for trust, and privacy rules that limit what a practice can measure. Plan around both from the start.

From ranking to being cited

For two decades, local search for a medical practice followed a predictable logic. A patient typed a condition and a city, a page of results appeared, and the practice that ranked well captured a share of clicks roughly proportional to its position. Strategy meant improving position.

That logic still applies, but it now sits underneath a second layer. Google's AI Overviews and AI Mode, along with assistants such as ChatGPT and Perplexity, increasingly answer the question before the patient sees a list. The answer summarizes what a condition is, what treatment usually involves and, often, which local providers are worth considering. A patient who reads that answer may contact a practice without ever scrolling to the organic results.1

The strategic consequence is a change in what is scarce. A results page has ten positions. A generated answer typically names a handful of sources and a few providers. The competition is narrower, the reward for being included is larger, and the penalty for being absent is close to total. For a practice owner, the question is no longer "where do we rank?" but "when a patient asks, are we part of the answer?"

Why healthcare is a special case

Every industry is adjusting to AI search, but healthcare carries two constraints that change the strategy rather than just the tactics.

The first is trust. Search systems have long applied a higher bar to health information, because a wrong answer can cause real harm.2 A generated answer about a medical procedure tends to lean on sources that look clinically credible: content with clear authorship, specific and accurate detail, and consistency with what reputable institutions say. A practice whose website reads like generic marketing copy gives the system little reason to cite it, however well it once ranked.

The second is privacy. Most industries can track every visitor, retarget them across platforms and feed that data into advertising algorithms without much thought. A medical practice cannot treat website behavior that way, because information collected on a provider's site can, in some circumstances, be protected health information. That limits the measurement playbook other industries take for granted, and any strategy that depends on it is built on sand.

A four-question framework

Building an AI in healthcare SEO strategy is less about new tools than about answering four questions in order. Each one depends on the one before it, and skipping ahead is the most common way practices waste effort.

1. Eligibility: can the systems read us at all?

Before a practice can be cited, its content has to be accessible to the systems doing the citing. Google's own documentation on AI features is notably plain on this point: the fundamentals that matter for search matter here too. Pages need to be crawlable and indexable, important information needs to exist as text rather than locked inside images, internal links need to make content findable, and any structured data needs to match what is actually visible on the page. Google has also said there is no special markup or separate file required to appear in these features.

This is the least glamorous question and the one practices most often skip. A well-written page that sits behind a misconfigured robots directive, or a list of services that exists only inside a graphic, will not be cited regardless of quality. Eligibility is a gate, not a differentiator, but nothing else matters until it is open.

2. Authority: why would an answer trust us?

Once a page can be read, the question becomes whether it deserves to be quoted. In healthcare, authority comes from specificity and accountability. A page explaining what to expect from a procedure should be written or reviewed by someone qualified to say it, should name that person, and should answer the questions patients actually ask in the order they ask them.

Generic pages lose this contest. A line such as this gives an answer engine nothing to use:

> We offer comprehensive care tailored to your needs

A clear explanation of who is a candidate for a treatment, what recovery usually involves and when a patient should seek a second opinion gives it something worth summarizing, and gives the practice a reason to be named. The strategic implication is uncomfortable for many practices: the content that earns citations is the content that takes clinical time to produce. It cannot be fully delegated to a copywriter with no medical input.

3. Consistency: does the rest of the web agree with us?

Generated answers do not rely on a practice's website alone. They draw on directory listings, business profiles, review platforms and any other source that describes the practice. When those sources disagree about the address, the specialties offered or the providers on staff, the system has less confidence in any of them, and a less confident system is less likely to name the practice at all.

That makes consistency a strategic asset rather than an administrative chore. A practice that describes itself the same way everywhere, keeps its physician profiles current and answers reviews in a measured, privacy-respecting way presents a coherent entity that an answer engine can describe with confidence. A practice with three phone numbers and an outdated provider list presents a puzzle, and puzzles get skipped.

4. Measurement: how will we know it is working?

The last question is where healthcare diverges most sharply from other industries. Traditional search reporting counts clicks and rankings. In an environment where a patient may read an answer and call without ever visiting the site, clicks understate impact, and rankings describe only part of the picture. Google has begun adding reporting on appearances in its AI features to Search Console, which helps, but the outcome that matters is still booked appointments.3

Measuring that outcome has to be done in a HIPAA-conscious way. The HHS guidance on online tracking technologies is a sensible starting point for any practice reviewing what its website collects and where that data goes. The practical answer for most practices is to measure at the front desk and the phone system, asking how new patients found the practice and tying call and form data to booked visits, rather than relying on third-party scripts that may expose patient information.

Where practices go wrong

Three strategic errors account for most of the wasted effort.

The first is treating AI search as a separate channel with its own budget and its own vendor. These systems draw on the same signals that drive organic search, so splitting the work splits the strategy. A practice that funds "AI optimization" separately from its website and listings often ends up paying twice for work that should have been done once.

The second is chasing shortcuts. Every shift in search produces a wave of tactics promising them: special files, keyword formulas, mass-produced pages aimed at every conceivable question. In a field where the systems are explicitly looking for trustworthy, specific health information, volume without substance is more likely to dilute a practice's credibility than to build it.

The third is measuring the wrong thing. A practice that judges its search strategy by traffic alone may conclude it is failing at exactly the moment it is succeeding, because a patient who gets a clear answer naming the practice may never need to click.

Sequencing the work

For a practice starting from a typical position, the four questions suggest a natural order. Fix eligibility first, because it is usually quick and everything else depends on it. Then rebuild the handful of pages that matter most to the practice's revenue, with clinician input and named authorship, rather than producing a large volume of new content. Clean up listings and profiles in parallel, since that work is largely administrative and can run alongside content. Put measurement in place early, so the practice has a baseline to judge against rather than a guess.

None of this produces results overnight, and any plan that promises fixed outcomes on a fixed timeline deserves caution. What it produces is a position: a practice that search systems can read, have reason to trust and can describe accurately.

Strategic lessons beyond healthcare

Healthcare is an early and demanding test of AI search, but the lessons travel. In any field where trust matters, from legal services to financial advice, the shift from ranking to citation rewards the same things: accessible information, accountable expertise, a consistent public identity and measurement tied to real outcomes rather than proxies. The organizations that do well will not be the ones that found the cleverest workaround. They will be the ones whose information was simply the most reliable thing available when the question was asked.

- 1Google users are less likely to click on links when an AI summary appears in the results

- 2KFF Health Misinformation Tracking Poll: Artificial Intelligence and Health Information
- 3Google adds AI Mode traffic to Search Console reports

Summary

Search is moving from a list of links to a synthesized answer, and the prize now goes to the practice named in that answer. Healthcare raises the bar through a higher trust standard and privacy rules that limit measurement. The four questions give a sequence: make content readable, earn authority through specific, accountable clinical content, keep every public description consistent and measure the outcome that matters, which is booked appointments. Practices that treat AI search as a separate channel, chase shortcuts or judge success by traffic alone risk wasting a year. The same logic extends to legal services and financial advice, where the most reliable information wins the citation. The advantage goes to the practice whose information is most dependable when the patient asks.