

# AI Search Readiness Assessment

perthharbourphysio.com.au

Prepared for Perth Harbour Physio    Location Fremantle, WA    Date April 2026    Readiness score 3 / 10

## Summary

This report looks at how well perthharbourphysio.com.au is set up to appear in AI-powered answers. That covers Google's AI Overviews, ChatGPT, Perplexity, Claude, and Gemini. If a parent asks ChatGPT "what's a good women's health physio in Fremantle?", the question is whether this clinic shows up in the answer or gets ignored.

Right now, the answer is ignored. I scored the site 3 out of 10. Three points because the site is fast, crawlable, and has 22 blog posts (which is more than most clinic websites I audit). The other seven points are missing because there is no structured data, no llms.txt, no FAQ sections, and the blog posts are short and inconsistent in topic coverage.

This matters more for a physio than for many other local businesses. Patients researching a specific condition (plantar fasciitis, pelvic floor issues, post-ACL rehab) increasingly ask AI first and check Google second. If the site isn't in the AI's training data or reachable via its citation tools, it never enters that conversation. The clinic's four physios have enough domain expertise to produce genuinely useful content. The site just isn't packaging that expertise in a format AI systems can use.

## Readiness score

**3 / 10**

OVERALL READINESS

**0%**

STRUCTURED DATA  
COVERAGE

**22**

BLOG POSTS, THIN  
& INCONSISTENT

**100%**

CRAWL ACCESS FOR  
AI BOTS

## How the score breaks down

Scoring is weighted toward signals AI systems actually use to identify and cite a business. A clinic scoring 9 or 10 would have Physiotherapist and Person schema everywhere, a structured blog with 30+ longer posts covering specific conditions, FAQ markup, llms.txt, and strong alt text coverage.

SIGNAL	STATUS	WEIGHT	SCORE
Site crawlable by all major AI bots	Yes	1.0	1.0
Sitemap present and valid	Yes	1.0	1.0
JSON-LD schema (LocalBusiness, Person, etc.)	None	2.0	0.0
llms.txt available at site root	404	0.75	0.0
FAQ sections and FAQPage schema	None	1.25	0.0
Educational content covering specific conditions	Partial (22 short posts)	2.0	0.7
Structured service and team information	Partial	1.0	0.3
Image alt text coverage	40%	0.5	0.2

## Detailed findings

### FIX FIRST

#### No schema markup anywhere on the site

For a physiotherapy clinic this is the biggest gap. Physiotherapist schema (a subtype of LocalBusiness) is exactly the signal Google, Bing, and AI citation tools look for when answering a query like "physio Fremantle". Without it, the clinic is indistinguishable from a blog with a contact page.

The markup AI search cares about most for this site:

- Physiotherapist on the homepage with address, phone, opening hours, payment accepted (HICAPS, Medicare, NDIS, DVA, WorkCover), and aggregateRating pulled from Google reviews
- Person on each practitioner bio with jobTitle, alumniOf (university), affiliation (APA membership where applicable), and knowsAbout (areas of clinical focus)
- Service on each service page, referencing the parent Physiotherapist
- FAQPage on service pages once FAQ blocks are added
- Article on blog posts with author attribution to the relevant physio

Author attribution in particular matters more than it used to. AI systems increasingly weight content by the demonstrable expertise of the author. Linking every blog post to a Person schema entry for the authoring physio, with their qualifications, is the cheapest credibility signal available.

**Time:** Half a day for a developer, or two hours with Rank Math / Schema Pro and a content pass.

## FIX FIRST

### Blog content is short and topic-scattered

22 posts is a good start. The problem is depth and focus. The average post is 420 words. The topics jump around: running injuries, an office ergonomics tip, a pelvic floor post from 2022, a generic "5 stretches for back pain" piece, a clinic Christmas update. Nothing goes deep enough to become the canonical source on a specific question.

AI systems cite the piece that answers the question best. That's usually the longer, more specific, better-structured piece. Short generic posts rarely get chosen because there is always a more thorough option.

A realistic plan for the next six months: write eight to ten pieces at 1,200 to 1,800 words each, each one targeting a specific question patients already ask on calls. Examples that would fit this clinic:

- "Plantar fasciitis: what we see in the clinic, what actually helps, and when to stop ignoring it"
- "Post-ACL reconstruction rehab: the timeline nobody tells you about"
- "Pelvic floor physio after a vaginal birth: what to expect in week one, week six, and month six"
- "NDIS physio in Perth: what plan-managed participants need to know about choosing a provider"
- "Returning to running after a calf strain: the five checkpoints we use with patients"
- "Workers' compensation physio in WA: your rights, the process, and what 'suitable duties' actually means"

Each of these answers a specific question patients are already asking. Each one is substantial enough to become the preferred citation. Each one positions a specific physio as the author, which feeds back into Person schema.

## FIX FIRST

### No FAQ content on any service or funding page

FAQs are one of the content formats AI systems pull from most reliably, because the question-and-answer structure mirrors how people actually prompt chatbots. For a clinic, the obvious FAQs are also the calls the receptionist already answers five times a day:

- Do you bulk bill? Do I need a GP referral to claim Medicare?
- Do you take NDIS plan-managed and self-managed clients? What about NDIA-managed?
- Can I claim this on my private health?
- Do you have a female physio available for women's health appointments?
- Do you offer WorkCover physio? How does that billing work?
- What's the difference between sports physio and general musculoskeletal physio?
- Do you do home visits? NDIS home visits?
- How soon can I get an appointment?

Put five to eight FAQs on the homepage, on each service page, and on the rebates/funding page. Wrap them in FAQPage JSON-LD. This single change will improve rankings for long-tail queries, improve AI citation odds, and reduce front-desk call volume. Three outcomes from one piece of work.

#### WORTH DOING SOON

### No llms.txt file

llms.txt is a plain-text file at the root of a site describing the business for large language models. It's still emerging as a standard, but Anthropic and a handful of others fetch it, and the cost to publish is 30 minutes.

For this clinic, the llms.txt should cover: the clinic name and location, the four practitioners and their specialties, the services offered, the funding types accepted, and links to the main pages. JimmyTech runs one live at jimmytech.com.au/llms.txt as a reference.

#### WORTH DOING SOON

### Service and team pages lack structured facts

The Sports Physio page describes sports physio in flowing prose. An AI reading it gets the gist but can't easily extract which conditions are treated, which techniques are used, which patient types are suitable, or who the practitioners are. The same is true on the other eight service pages.

Convert the key facts to lists and small tables. For each service page:

- A "Who this helps" bullet list of conditions
- A "Typical treatment approach" list of techniques
- A "Practitioners who offer this" section linking to Person bios
- A "Funding accepted" mini-list (Medicare EPC, WorkCover, NDIS, private health, DVA)
- A small FAQ block

None of this requires replacing the current prose. It's additive. The prose still reads well to a human; the lists give the AI something clean to extract.

#### WORTH DOING SOON

### Image alt text at 40%

67 of 112 images have no alt text. Multimodal AI (Google AI Overviews, GPT-4 Vision, Claude) can read images alongside the text on a page. Practitioner headshots, clinic interior photos, and treatment demonstration photos all carry real signal. Without alt text, that signal is lost.

For a physio clinic, good alt text is specific to what's in the image. "Mel Harrison demonstrating pelvic floor activation with a patient seated on a fitball" is useful. "Blog header" is not. Practitioner headshot alt text should include the physio's name and title.

#### WORTH DOING SOON

### Author bylines are weak

Blog posts are attributed to "Admin" or to a generic "Perth Harbour Physio" account. No links to practitioner bios, no author schema, no qualifications listed under the byline. Posts feel anonymous, which weakens their credibility signal for both Google and AI.

Every post should be bylined by the specific physio who wrote it, linked to their bio page, and carry an Article schema with an `author` field pointing at the Person schema on the bio. That's the full trust chain AI systems follow.

#### NICE TO HAVE

### robots.txt is the default file

All AI crawlers are allowed through by default, which is the right outcome for a clinic that wants visibility. It's working by accident rather than by design. Add explicit `User-agent` blocks for GPTBot, ClaudeBot, Google-Extended, PerplexityBot, and CCBot with `Allow: /`. Adds intention to the policy.

#### WORKING

### Site is fully accessible to AI crawlers

No blocks on any major AI bot. GPTBot, ClaudeBot, Google-Extended, ChatGPT-User, PerplexityBot, CCBot all get through. Sitemap is valid. This is the right starting point.

## Content structure, page by page

AI systems extract answers more reliably from structured content: clear headings, lists, tables, FAQ blocks. Here's how the key pages rate for machine readability.

### Homepage / · Moderate

Decent structure. H1, four H2 sections (About, Services, Team, Book Online). The hero carries the clinic name and a tagline. Some key facts sit in bullet form (insurance accepted, location).

Gaps: no schema, no FAQ block, no clear answer to "what does this clinic specialise in?" for a reader in a hurry. An AI summarising this page would produce something generic: "a physiotherapy clinic in Fremantle". It wouldn't know about sports, women's health, or NDIS as specialisations.

## **Sports Physiotherapy** /services/sports-physiotherapy/ · Moderate

The strongest of the service pages. Around 650 words. Mentions specific conditions (ACL, calf strains, shoulder impingement). Lists the practitioners who offer this service.

Gaps: no FAQ block, no Service schema, no structured list of conditions treated, no clear funding information.

## **Dry Needling** /services/dry-needling/ · Low

112 words. Two short paragraphs. No structure beyond the title. An AI asked "does Perth Harbour Physio offer dry needling?" could answer yes, but couldn't say what it's used for, which physios offer it, what a session costs, or how it's different from acupuncture.

## **NDIS Physiotherapy** /services/ndis-physiotherapy/ · Low to moderate

340 words, slightly longer than Dry Needling but still thin for a page that should be the clinic's primary entry point for NDIS participants. It mentions that the clinic accepts plan-managed and self-managed participants. It doesn't cover home visits, what kinds of NDIS goals the physios work toward, what NDIS reports look like, or how to book as a plan-managed participant.

This page has the highest upside of all the service pages. NDIS search traffic in WA is growing and the competition writes about it poorly. Doubling this page's length with genuinely useful information for support coordinators and plan managers would pay back within six months.

## **Team bios** /team/\* · Low to moderate

Each of the four practitioner bios is 180 to 280 words. They include qualifications and interests. They do not include Person schema, links to services they offer, links to blog posts they've written, or any structured facts about experience, registrations, and affiliations.

Person schema with `jobTitle`, `alumniOf`, `hasOccupation`, and `knowsAbout` turns each bio into a credible entity AI systems can cite. Without it, the physios exist on the site but not in the knowledge graph.

## **Blog posts** /blog/\* · Varies

22 posts, average 420 words. The longer ones (two 900-word posts on running injuries and pelvic floor recovery) are the only pieces on the site with enough depth for an AI to cite confidently. The rest are too short to win against better-written alternatives on the same topics.

# Prioritised recommendations

## 1. Add schema markup across the site

Start with Physiotherapist on the homepage, Person on each team bio, Service on each service page, and BreadcrumbList sitewide. Layer Article and FAQPage in as blog and FAQ content expands. This is the single highest-impact change per hour of work and has knock-on benefits for Google rich results, Google Business Profile, and AI citation.

**Time:** Two to four hours with a schema plugin. Validate with Google's Rich Results Test before deploying.

## 2. Publish a series of deep condition-specific articles

One in-depth article per month for the next six months. 1,200 to 1,800 words each. Each one answers a real patient question. Each one authored by the physio who actually treats that condition. This is how AI search starts pointing patients at the clinic.

The first six topics should be decided together. My suggestions above are a starting point; the team will have better instincts about what patients actually ask and what's under-covered in existing search results.

## 3. Add FAQ sections with FAQPage schema

Five to eight FAQs on the homepage, each service page, and the rebates page. Write answers the way the front desk answers them on calls. Wrap in FAQPage JSON-LD.

**Time:** Two to three hours to write, 30 minutes for the schema.

## 4. Publish an llms.txt file

Write a structured file at [perthharbourphysio.com.au/llms.txt](https://perthharbourphysio.com.au/llms.txt) describing the clinic, practitioners, services, funding accepted, and the main pages. Thirty minutes of work.

## 5. Expand the NDIS and funding pages

Rewrite the NDIS page to 900+ words covering home visits, report types, typical goals, how billing works, and practical guidance for plan managers and support coordinators. Do similar on the WorkCover and Medicare EPC pages. These pages sit at the top of the funnel for high-intent visitors.

## 6. Restructure service pages with lists and tables

Keep the existing prose. Add structured blocks: conditions treated, techniques used, practitioners, funding accepted, FAQ. This makes the pages more scannable for humans and more extractable for AI.

## 7. Improve author bylines and attribution

Every blog post should be bylined to a specific physio, linked to their bio page, and carry Article schema with the correct author reference. This is a credibility chain AI systems follow.

## 8. Fix image alt text across the site

Write specific, descriptive alt text for the 67 images missing it. Use practitioner names on headshots. Describe what's happening in treatment photos.

# What "good" looks like for a physio website

A clinic that's well-set-up for AI search has most of this in place:

- Physiotherapist schema on the homepage with full NAP, hours, payment types, and aggregateRating
- Person schema on each practitioner page with qualifications, affiliations, and areas of expertise
- Service schema on each service page
- FAQ blocks on every service page and the funding page, wrapped in FAQPage schema
- A blog averaging 1,200+ words per post, focused on specific conditions, authored by the relevant physio
- Article schema with author references linking back to Person schema
- llms.txt at the site root
- Descriptive alt text on every clinical or team photo
- An explicit robots.txt policy for AI crawlers

---

Prepared by **Jimmy Faccioli** — JimmyTech, Perth WA

This is a sample report for a fictional clinic. Your audit will be tailored to your site's actual findings. [jimmytech.com.au/seo-audit](https://jimmytech.com.au/seo-audit)