

● SITE TEARDOWN · DOCTORGPEDS

Doctorgpeds: a phone-and-desktop teardown

We loaded <https://doctorgpeds.com/> 4 times on a simulated iPhone and 2 more on a 1366px desktop, and wrote down what a real visitor would see on each. No login, no insider access, no Harvv pixel needed. Here is what repeated visits already show, sorted by how we know it.

August 8, 2026 · External scan · 4 MOBILE + 2 DESKTOP LOADS · NO PIXEL DATA

TL;DR

What jumped out

Of everything we found on this scan, this is the one to start with: **This page is telling Google not to list it.** The page carries a robots "noindex" tag, which asks search engines to leave it out of results entirely. If this page is meant to be found, that tag is silently costing you every bit of organic traffic. Often a leftover from staging.

18%

TAP TARGETS UNDER THE
44PX MINIMUM (3 OF 17)

172ms

LARGEST CONTENTFUL PAINT
(LAB MEDIAN)

0.2 MB

TOTAL PAGE WEIGHT (MEDIAN)

Below: what's already working, every finding ranked by impact and tagged with the screen it affects, the speed numbers on phone and desktop, and a checklist of what to fix first.

AI SEARCH VISIBILITY

How visible is doctorgpeds.com to AI search?

ChatGPT, Perplexity, and Google AI Overviews are the new front page. The AI checks in this report are a preview. A free account turns them into a live 0 to 100 score, tracked daily and benchmarked against your competitors.

[See the full AI report, free](#)

00

What's already working

Start here so the problems below are in context. These held up across the test loads:

- **Speed is good.** Google Lighthouse scores this page 92/100 for performance on mobile. Loading is not what is costing you visitors here.
- **Light page weight.** The page is about 0.2 MB across 10 requests. That keeps it quick on mobile data and cheap to load repeatedly.
- **No JavaScript errors on load.** Nothing threw a script error across the test loads, so buttons and tracking are not silently breaking mid-session.
- **Accessible to most visitors.** Lighthouse accessibility is 100/100, so screen-reader and contrast basics are largely handled.
- **Layout holds on phone and desktop.** Nothing spilled past the edge at either 390px (phone) or 1366px (desktop), so the structure is responsive.

01

Findings, ranked by what hurts conversion most

SEVERITY	FINDING	HOW WE KNOW
HIGH	This page is telling Google not to list it The page carries a robots "noindex" tag, which asks search engines to leave it out of results entirely. If this page is meant to be found, that tag is silently costing you every bit of organic traffic. Often a leftover from staging.	BOTH

SEVERITY	FINDING	HOW WE KNOW
HIGH	Add quotations so AI engines cite this page BOTH AI SEARCH SEO Generative engines (ChatGPT, Perplexity, AI Overviews) lift sourced, attributed quotes almost verbatim, and quotations are the single strongest citation lever (studies measure roughly +41%). Add 1-2 attributed expert quotes or blockquotes to the pages below. (Found across a sample of 5 pages from your sitemap, a partial crawl rather than your full site.)	
HIGH	Broken internal links BOTH SEO CONVERSION 1 EXAMPLE 1 internal link target returned an error (404/410/5xx) or did not respond. Broken links waste crawl budget, frustrate visitors, and leak ranking signal. (Found across a sample of 5 pages from your sitemap, a partial crawl rather than your full site.) ----- THE EXACT EXAMPLES WE FOUND: https://doctorgpeds.com/cdn-cgi/l/email-protection (404)	
MEDIUM	AI search crawlers are blocked BOTH SEO 8 EXAMPLES Your robots.txt blocks 8 AI crawlers (gptbot, google-extended, claudebot, ccbot, bytespider, applebot-extended, amazonbot, meta-externalagent). Those engines (ChatGPT, Perplexity, Google AI Overviews, Claude) cannot read your site, so your content cannot be cited or shown in AI search. ----- THE EXACT EXAMPLES WE FOUND: - gptbot - google-extended - claudobot - ccbot - bytespider - applebot-extended - amazonbot - meta-externalagent	
LOW	Security headers are missing or weak BOTH SECURITY The server response is missing browser-hardening headers that protect visitors and are a standard security and agency checklist item. Missing or weak here: HSTS (Strict-Transport-Security, which forces HTTPS on return visits). These are set at the server, CDN, or host level (most platforms expose them in settings or a config file) and do not change how the site looks or performs.	

SEVERITY	FINDING	HOW WE KNOW
LOW	No third-party analytics tag detected BOTHTRACKING No Google Analytics, GA4, or other third-party analytics tag was found in the page. If you rely only on platform or server-side analytics that is fine; but if not, a third-party tag is the baseline way to see how many visitors arrive, where they come from, and what converts. First-party or server-side analytics are not detectable from the outside, so this may already be handled.	
LOW	No email capture or newsletter detected BOTHCONVERSION No email-marketing tag (Klaviyo, Mailchimp, etc.) was found. If capturing visitors is relevant to this site, an email signup builds an owned audience you keep regardless of ad costs. Not every site needs one.	
LOW	Sitemap not listed in robots.txt BOTH A sitemap exists but is not referenced in robots.txt. Adding a "Sitemap:" line to robots.txt helps crawlers find it reliably.	
LOW	No Organization or WebSite schema BOTHSEO No site-wide Organization or WebSite structured data found on any crawled page. This is the schema that tells Google and AI engines who you are (name, logo, social profiles) for the knowledge panel and brand recognition in AI answers. (Found across a sample of 5 pages from your sitemap, a partial crawl rather than your full site.)	

FROM FINDING TO FIX

Want the fix, not just the finding?

Install Harvv and we turn each issue above into a ready-to-paste prompt for your AI coding assistant. Drop it into Cursor, Claude, or Copilot and the diagnosis becomes a concrete code change, written against this exact page.

Get the fix prompts

"How we know": **unlabeled** = a deterministic fact, identical on every load (e.g. element sizes). Most findings are this kind, so we only mark the exceptions: **median across loads** = a noisy lab metric, reported as a median. **real-user field data** = Google CrUX, actual Chrome visitors.

Structural and AI-search checks crawl up to 8 pages from your sitemap (a sample, not your full site). "Broken" means a link returned 404, 410, or 5xx, or did not respond; access-controlled pages (401, 403) are not counted.

02

Performance: phone, desktop, and real visitors

METRIC	MOBILE	DESKTOP	READ
TTFB (lab median)	27 ms	23 ms	LAB
FCP (lab median)	168 ms	168 ms	LAB
LCP (lab median)	172 ms	168 ms	GOOD
Page weight (median)	0.2 MB	0.2 MB	OK

Google Lighthouse (lab): Performance **92** mobile / **100** desktop, SEO **69**, Accessibility **100**, Best Practices **100**.

Lab numbers are from a headless mobile browser on an unthrottled connection: treat them as a floor, not a typical experience. Add a Google API key to light up real-user field data (CrUX) and Lighthouse scores.

03

Tiny buttons are hard to tap on mobile

3 of 17 tappable items on this page come in below 44×44 pixels, the size Apple and Google both recommend for reliable tapping on a phone. The same ones came up small on every one of the 4 test loads, so this is the page itself, not a fluke.

When customers can't tap what they expect to, they get frustrated and many of them leave. They don't file a bug. They don't try again. They just leave. A desktop dashboard can't see this because it's the difference between a thumb and a cursor.

The buttons measuring below the minimum on this scan:

- a 87x20 "Contact us"
- a 196x24 "Notice of Privacy Practices"
- a 103x24 "Privacy Policy"

The fix is CSS-only on most sites: add padding around the icon (don't just change the icon size) so the actual tap area is at least 44×44 pixels. No redesign, no new assets.

04

Technical SEO & structured data

CHECK	RESULT
Title	Dr. G Pediatrics Fresno Pediatrician & Children's Clinic (58 chars)
Meta description	126 chars
H1	1 on page
Canonical	PRESENT
Structured data (JSON-LD)	MedicalClinic
Open Graph	TITLE + IMAGE

05

The fix checklist

Everything to fix, priority first, each tagged with the screen it affects and a rough effort. Work top to bottom.

☐ This page is telling Google not to list it BOTH VARIES

☐ Add quotations so AI engines cite this page BOTH VARIES

☐ Broken internal links BOTH VARIES

☐ AI search crawlers are blocked BOTH VARIES

☐ Security headers are missing or weak BOTH VARIES

☐ No third-party analytics tag detected BOTH DEV AFTERNOON

☐ No email capture or newsletter detected BOTH VARIES

☐ Sitemap not listed in robots.txt BOTH VARIES

☐ No Organization or WebSite schema BOTH VARIES

Effort is a rough read from the outside: "CSS only" means no new assets or backend work, "1 line" means a single tag, "Dev afternoon" means a developer needs to touch tracking or scripts.

06

What this report cannot tell you

Everything above is from the outside, looking at the page on a simulated phone and desktop. The questions that actually decide revenue need real visitors. Install the Harvv pixel (one script tag, 16 KB, zero personal data, no engineering project) and within about 72 hours you'd know which buttons real customers tapped and missed, how often Google Analytics is missing visits, and exactly where mobile shoppers stalled and left. This report shows you where to look. The pixel shows you how often it happens, and to whom.

07

How we did this, and what it can't prove

- **4 mobile + 2 desktop loads** of one URL from headless Chrome (iPhone viewport at 390px, desktop at 1366px), August 8, 2026. Enough loads to separate real defects from random noise, not a full-site crawl.
- **Lab numbers, not real-user numbers** (no field data was available for this run). Real devices on real networks run slower.
- **Friction is inferred, not counted.** We can prove a button is small. We can't, from the outside, count how often it causes a missed tap. That requires the pixel on a live page.

About Harvv, the source of this teardown

Harvv is a behavioral UX analytics platform (harvv.com). A lightweight JavaScript pixel captures how real visitors behave on a site (dead clicks, rage clicks, scroll depth, Core Web Vitals, JavaScript errors, and 50+ other signals) and the engine turns them into prioritized, plain-English findings. This teardown is the outside-in version of that: the same detectors run against a public page, with no pixel installed.

How to read it. Every finding here is a reproducible, automated measurement, not an opinion: element sizes, contrast ratios, load metrics, and structured-data checks that anyone can re-run against the same URL. The

method is stated in full above. Automated testing catches a subset of issues, so this is a starting point, not a certification.

Full disclosure. Harvv makes the pixel that would measure the friction these findings imply, so we have a commercial interest. That is exactly why the findings are kept to things a reader can verify independently, and why nothing here is inflated: an unreproducible claim would undermine the tool it is meant to demonstrate.

Prepared by Harvv (harvv.com), a behavioral UX analytics platform. Last updated August 8, 2026.