

● SITE TEARDOWN · MIHIRNAIK

Mihirnaik: a phone-and-desktop teardown

We loaded <https://mihirnaik.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 18, 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: **Some text is low-contrast and hard to read.** Text that does not stand out enough from its background is hard to read for many visitors, and fails accessibility guidelines Google checks.

18%

TAP TARGETS UNDER THE
44PX MINIMUM (12 OF 65)

984ms

LARGEST CONTENTFUL PAINT
(LAB MEDIAN)

0.0 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 mihirnaik.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.** The main content paints in about 1.0s in our test loads, inside Google's 2.5s "good" threshold. Real networks are slower, but the page itself is not heavy.
- **Light page weight.** The page is about 0.0 MB across 97 requests. That keeps it quick on mobile data and cheap to load repeatedly.
- **Search basics are in place.** Lighthouse scores SEO 100/100. The fundamentals Google looks for are present.
- **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	Some text is low-contrast and hard to read Text that does not stand out enough from its background is hard to read for many visitors, and fails accessibility guidelines Google checks.	BOTHACCESSIBILITY (WCAG)
HIGH	JavaScript errors were logged during load Errors in the console often mean a feature silently broke: a button that does nothing, analytics that did not fire, a form that will not submit.	BOTHCONVERSIONTRACKING

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 8 pages from your sitemap, a partial crawl rather than your full site.)	
MEDIUM	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: Content-Security-Policy (the main defense against injected and cross-site scripts); clickjacking protection (X-Frame-Options or a CSP frame-ancestors rule); HSTS (Strict-Transport-Security, which forces HTTPS on return visits); X-Content-Type-Options: nosniff (stops MIME-type sniffing attacks). 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.	
MEDIUM	5 potential dead-click targets BOTH CONVERSION ACCESSIBILITY (WCAG) TRACKING 5 ELEMENTS Elements styled like buttons but with no anchor, no <button> wrapper, no role="button", and no click attribute. Real visitors tap these expecting something to happen, then leave. Examples on this page: "Resources" (span.header-nav-link), "Services" (span.header-nav-link), "Projects" (span.header-nav-link). ----- THE EXACT ELEMENTS WE FOUND: "Resources" (span.header-nav-link) "Services" (span.header-nav-link) "Projects" (span.header-nav-link) "Perplexity" (div.hero-rotator-pill) "Share your challenge and we'll see if we're a fi" (div.ant-typography)	

SEVERITY	FINDING	HOW WE KNOW
MEDIUM	5 interactive elements have no stable, accessible identity ACCESSIBILITY (WCAG) TRACKING 5 ELEMENTS	BOTH
	These elements are clicked like buttons but expose no accessible name, or are a plain div/span used as a control with no role. Assistive tech announces only a role (or nothing), and analytics and heatmaps have no human-readable label or stable selector to bind the click to, so the click is both inaccessible and untrackable, and any redesign silently breaks click aggregation. Give each one a real <code><button>/<a></code>, an aria-label, and a stable id or data-attribute. THE EXACT ELEMENTS WE FOUND: • <u>"Resources" (span.header-nav-link) (div/span used as a button, no role)</u> • <u>"Services" (span.header-nav-link) (div/span used as a button, no role)</u> • <u>"Projects" (span.header-nav-link) (div/span used as a button, no role)</u> • <u>"Perplexity" (div.hero-rotator-pill) (div/span used as a button, no role)</u> • <u>"Share your challenge and we'll see if we" (div.ant-typography) (div/span used as a button, no role)</u>	
LOW	Images have no width or height set (layout is stable for now) PERFORMANCE SEO 6 IMAGES	BOTH
	46 of 72 images don't declare width and height. Your layout currently stays stable (measured Cumulative Layout Shift is 0, which is fine), most likely because space is reserved another way (CSS), so this is not causing visible jumps today. It is still worth setting explicit dimensions or a CSS aspect-ratio so a slow connection or a future style change cannot reintroduce shift. THE EXACT IMAGES WE FOUND: • <u>perplexity.svg</u> • <u>airnewzealand.svg "Air New Zealand"</u> • <u>att.svg "AT&T"</u> • <u>bbc.svg "BBC"</u> • <u>capgemini.svg "Capgemini"</u> • <u>clickup.svg "ClickUp"</u>	

SEVERITY	FINDING	HOW WE KNOW
LOW	No canonical tag, so duplicate URLs split the page's ranking BOTH When the same content is reachable at multiple URLs (think tracking parameters or session IDs), Google can split your ranking signal across them. A single canonical tag tells Google which version counts.	
LOW	Search-result title will be cut off BOTH SEO The title is 74 characters; Google truncates around 60. Front-load the important words so nothing useful gets clipped.	
LOW	No third-party analytics tag detected BOTH TRACKING 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 BOTH CONVERSION 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	No visible contact details (email or phone) BOTH CONVERSION The page exposes no email or phone link. For higher-value or trust-sensitive purchases, a clear way to reach a human reduces hesitation. Add an email or phone link in the header or footer.	
LOW	Unused JavaScript is being downloaded BOTH PERFORMANCE Code that never runs on this page still costs download and parse time on every visit. Splitting or removing it speeds up load. Lighthouse measured: Est savings of 85 KiB.	
LOW	Unused CSS is being downloaded BOTH PERFORMANCE Style rules that this page never uses still block rendering while they download. Trimming them frees the paint path. Lighthouse measured: Est savings of 31 KiB.	
LOW	Headings skip levels BOTH ACCESSIBILITY (WCAG) Jumping from an H1 straight to an H4 breaks the document outline that screen readers and search engines rely on.	
LOW	No /llms.txt file BOTH SEO No /llms.txt. Optional: some AI tools read this file (Anthropic, Vercel, Stripe, Cloudflare and Hugging Face publish one, and AI coding assistants use it to find docs), so adding one is a cheap courtesy to them. Google has said in writing that Google Search and its AI features do not use llms.txt, so do not expect a search-ranking or AI Overviews benefit from it.	

Accessibility findings are automated checks against Web Content Accessibility Guidelines (WCAG) 2.1 and 2.2. They flag potential barriers and legal risk, not a certification or a determination of compliance with the ADA, Section 508, or EN 301 549. Automated testing catches only a subset of issues; a full conformance review needs manual and assistive-technology testing by a qualified reviewer.

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)	585 ms	594 ms	LAB
FCP (lab median)	984 ms	996 ms	LAB
LCP (lab median)	984 ms	996 ms	GOOD
Page weight (median)	0.0 MB	0.1 MB	OK

Google Lighthouse (lab): Performance **66** mobile / **66** desktop, SEO **100**, Accessibility **95**, Best Practices **96**.

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

12 of 65 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 115x20 "Read full bio →"
- a 94x20 "Clarity ArcAI"
- a 191x20 "consulting engagements"
- a 110x22 "LinkedIn 18k+"
- a 69x22 "X 2.8k+"
- a 120x22 "Newsletter 2k+"
- a 103x22 "YouTube ~1k"
- a 63x22 "GitHub"

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	Mihir Naik AI Search (AEO) Professional, Product Leader & Enterprise SEO (74 chars)
Meta description	246 chars
H1	1 on page
Canonical	MISSING
Structured data (JSON-LD)	2 block(s)

CHECK	RESULT
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.

☐ Some text is low-contrast and hard to read BOTH VARIES

☐ JavaScript errors were logged during load BOTH DEV AFTERNOON

☐ Add quotations so AI engines cite this page BOTH VARIES

☐ Security headers are missing or weak BOTH VARIES

☐ 5 potential dead-click targets BOTH CSS ONLY

☐ 5 interactive elements have no stable, accessible identity BOTH DEV AFTERNOON

☐ Images have no width or height set (layout is stable for now) BOTH CSS ONLY

☐ No canonical tag, so duplicate URLs split the page's ranking BOTH 1 LINE

☐ Search-result title will be cut off BOTH 1 LINE

☐ No third-party analytics tag detected BOTH DEV AFTERNOON

☐ No email capture or newsletter detected BOTH VARIES

☐ No visible contact details (email or phone) BOTH VARIES

☐ Unused JavaScript is being downloaded BOTH VARIES

☐ Unused CSS is being downloaded BOTH VARIES

☐ Headings skip levels BOTH VARIES

☐ No llms.txt file 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 18, 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 18, 2026.