

AI Diligence & Transparency Statement

Tai Ku Labs believes radical transparency is the only honest answer to “do you use AI?” This statement lays out exactly where AI shows up in how we build and operate, and where it doesn’t.

How We Use AI

We use AI tools, primarily Claude (Anthropic’s AI assistant), as a production collaborator in specific, bounded ways. We treat it the same way we treat any other tool in the studio: useful for the 80% of work that is repeatable, never in charge of the 20% that is creative, strategic, or relational.

Concretely, here is where AI has shown up in building Tai Ku Labs as a company:

- **Website development.** The tai-ku.com site was built collaboratively with Claude. AI assisted with React scaffolding, CSS and layout, responsive behavior, animation iteration, performance work (de-inlining images to hashed, cache-friendly assets, converting photography to modern formats, and code-splitting the page so it loads fast), accessibility work (keyboard navigation, screen-reader semantics, contrast, and reduced-motion support), the motion design system, individual project pages, and back-and-forth copy refinement under our direction.
- **Copy refinement.** Headlines, service descriptions, project blurbs, and the About section were drafted by hand, refined in dialog with AI, and approved line by line by David Hernandez. No copy ships without a human pass.
- **Data tooling.** Our Master CRM (built in Attio) uses AI-assisted scripts for contact deduplication, schema normalization, and batch imports from event guest lists and outreach spreadsheets. Contact records themselves are not sent to training systems.
- **Form wiring.** The website’s contact form and newsletter signup post to secure Cloudflare Pages Functions that we own. Those functions then relay to Attio (our Master CRM) and Beehiiv (for newsletter delivery). API keys are stored as encrypted environment variables on our server, never exposed to the browser. AI assisted in writing the proxy logic and validation code; the data stays in our own systems.
- **Operations support.** Ideation, meeting prep, research summaries, and drafting documents like this one.
- **Site monitoring.** Automated AI-run checks verify weekly that the site is healthy (content serving, forms reachable, security headers intact) and audit quarterly how accurately AI

systems describe Tai Ku Labs. These checks are read-only: they report to a human and never change anything themselves.

How We Use AI to Check AI

One thing we want to be specific about, because it is unusual and because it is the part we are most careful with: we use AI to audit AI, with a human holding the final call at every step.

The build runs as a loop between two roles. Claude in an orchestration and review seat plans the work, reasons about trade-offs, and checks results. Claude in an execution seat (a coding agent working directly in the codebase) carries out the changes, runs tests, and reports back. Each side reviews the other's work: the planner catches when the executor's output drifts from intent, and the executor surfaces real runtime evidence that corrects the planner's assumptions. When one of them was confidently wrong, the other was usually the thing that caught it.

That cross-checking is not a substitute for human judgment, it is a tool that makes human judgment faster. David Hernandez sits in the loop as the decision-maker. Nothing ships on an AI's say-so. Every gate, the contact form, the newsletter wiring, the security headers, the analytics, the deployment itself, was tested live by a human before it was trusted, and more than once a green "success" message on screen was rejected until the underlying record was confirmed in our own systems. The discipline is simple: AI can propose, AI can verify another AI, but a person tests it against reality and decides.

Built to Be Read by AI

More and more prospective clients ask an AI assistant to research or shortlist agencies. We think the honest response is to make sure those systems get accurate, verifiable information, so we build for machine readers on purpose:

- **The same content for machines and humans.** Every project, service, and fact on tai-ku.com is embedded in the page in plain, readable form, so AI crawlers that don't run code see exactly what a human visitor sees. We do not show machines anything different from what people see.
- **An open brief for AI agents.** We publish /llms.txt (a summary of who we are, what we do, and who we've worked with) and /llms-full.txt (our complete project archive) specifically for AI systems, following the emerging llms.txt convention.
- **Structured data.** The site carries machine-readable descriptions (schema.org) of our organization, services, flagship events, and frequently asked questions.

- **AI crawlers welcome.** Our robots.txt explicitly permits AI crawlers (GPTBot, ClaudeBot, PerplexityBot, and others) to index and cite us.
- **The standard still applies.** Everything published to machines meets the same bar as everything published to humans: factual, verifiable, and approved by a person. If an AI recommends us, we want it to be because the record earned it.

Site Infrastructure (updated July 2026)

For full transparency about how tai-ku.com is built and operated:

- **Hosting.** Cloudflare Pages, with Cloudflare serving DNS, CDN, and DDoS protection.
- **Source.** Single-page React application built with a modern bundler. Images and brand assets are served as separate, content-hashed files so browsers cache them efficiently, and below-the-fold sections load on demand to keep the first paint fast.
- **Forms.** Contact and newsletter signup go through Cloudflare Pages Functions (serverless functions on our own origin). These functions validate input, run a spam honeypot check, enforce an origin allowlist, and then forward to Attio and Beehiiv using credentials held only on the server.
- **Security.** The site ships standard protective response headers (HTTPS enforced via HSTS, clickjacking protection, content-type and referrer hardening, a permissions policy, and an enforced content security policy that allows no inline scripts at all). Our form endpoints sit behind a rate-limiting rule to blunt abuse and spam floods. Secrets live only as encrypted server-side environment variables.
- **Analytics.** Google Analytics 4, aggregate and anonymized only. No cross-site tracking, no advertising cookies.
- **Privacy policy.** Published at /privacy.html, accessible from the footer of every page.

What AI Does Not Do For Us

- **No AI-generated imagery.** Every photograph on tai-ku.com and in our work product is original, shot by David Hernandez or a photographer we've hired. No diffusion-model art, no synthetic portraits, no generated event crowd shots.
- **No AI-generated video.** All video on our site and in our client decks is captured by human operators.
- **No synthetic event coverage.** Attendance figures, press mentions, awards, and project details are factual and verifiable. If we say Clio, we have the trophy. If we say 2,500 people, we had 2,500 people.

- **No AI-written client deliverables without human authorship.** When we deliver strategy, decks, or creative direction to a client, a human has written and stands behind every page.
- **No training on client or community data.** We do not share partner, client, or CRM data with AI training pipelines.

What Remains Human-Led

- **Creative direction.** All visual, strategic, and conceptual decisions come from David Hernandez and the Tai Ku team.
- **Photography.** Every image is original, shot on film or digital by our photographers.
- **Event production.** Curation, design, and execution of our IRL experiences (from Web Flea to Vibe Check to brand activations) is guided entirely by human judgment and lived cultural experience.
- **Brand identity.** Tai Ku's visual language, tone, and strategic positioning are authored by humans.
- **Client relationships.** Partnerships with brands, communities, and collaborators are built on human connection.
- **Final judgment.** The call on whether something is on-brand, on-culture, or ready to ship is always made by a human with context.

Data & Privacy

We do not share customer data, contact information, or proprietary business data with AI training systems. When we use AI tools internally:

- We follow Anthropic's privacy policy and enterprise data handling standards, including opt-out of model training on our inputs.
- Client information, event data, and brand assets are treated with the same confidentiality offline and online.
- We retain full ownership of our creative output, strategy, and intellectual property.
- Analytics on tai-ku.com are limited to Google Analytics 4 aggregate traffic data. We do not sell visitor data.
- Newsletter signups are stored in our own Attio CRM and Beehiiv account. We do not rent or sell our list.

Why This Matters

AI is a tool in our kit, like Lightroom, like Canva, like a project tracker. It doesn't replace what makes Tai Ku Labs worth hiring: cultural insight, an eye for design, the ability to feel a room, and the track record of actually pulling off disruptive events in cannabis, Web3, wrestling, food, and culture.

We're not hiding behind AI. We're not pretending it doesn't exist. We use it where it earns its place: compressing repetitive work, supporting ideation, speeding up technical scaffolding, and even checking its own work, while keeping creative ownership and final judgment exactly where they belong, with us and our collaborators.

Ongoing Commitment

As AI tools evolve, our practice will evolve with them. We commit to:

- Maintaining transparency about our use of AI in our work.
- Respecting the creative and intellectual contributions of our team and collaborators.
- Prioritizing human experience and cultural authenticity over algorithmic efficiency.
- Revisiting and updating this statement as tools, norms, and our practices change.

Questions or concerns? Reach out: info@tai-ku.com

Statement last updated: July 2026