

JAMESKICKLIGHTER.COM

The Custom Systems Handbook

Every custom-built tool and visibility system on the site — in plain English. How the films, photos, and writing get found by Google, by AI assistants, by voice, and everywhere in between.

3

SURFACES OF SEARCH

70+

CUSTOM TOOLS

1

CONNECTED SYSTEM

Prepared June 2026 · A reference for the owner and collaborators · No technical knowledge required

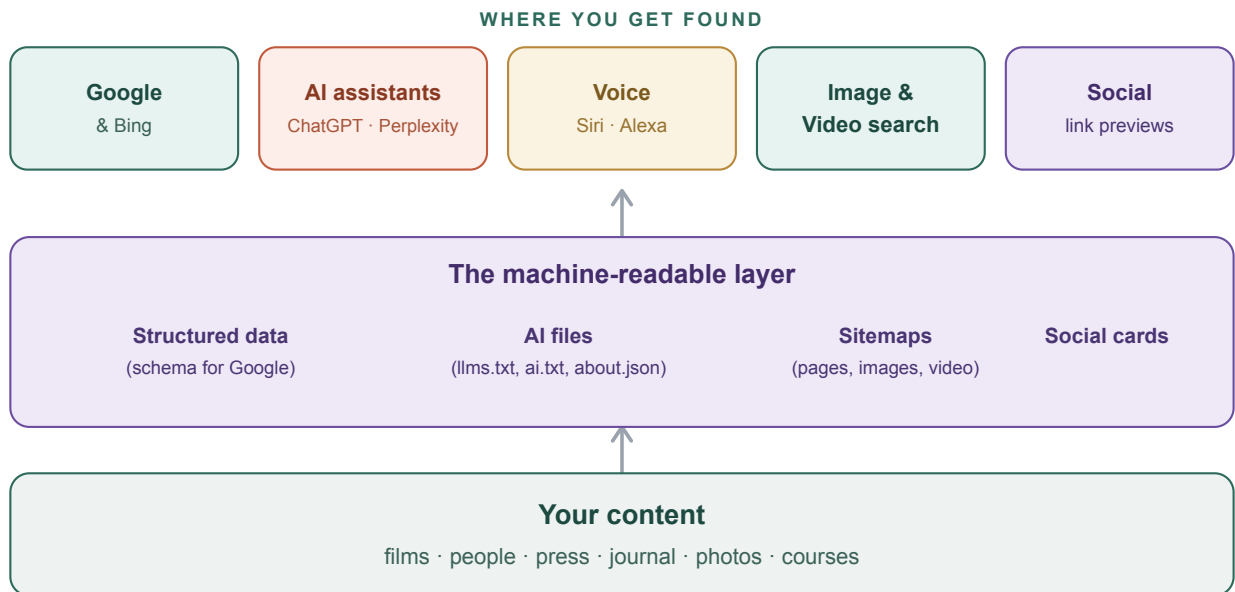
START HERE

The idea in one page

The website isn't just a set of pages. It's a machine that takes everything you create — films, people, press, journal entries, behind-the-scenes photos — and makes it findable wherever someone (or something) goes looking: Google, Bing, ChatGPT, Perplexity, Claude, Gemini, voice assistants, image search, and social feeds.

Most websites are built only for human eyes. This one is built for **both humans and machines**. Every page quietly publishes a second, machine-readable version of itself — a clean summary a search engine or an AI can read without guessing. That's the difference between "having a website" and "being an entity the whole internet understands."

There are three layers, and the whole handbook is organized around them:



The whole system, bottom to top: your content is translated into a machine-readable layer, which makes it visible across every kind of search.

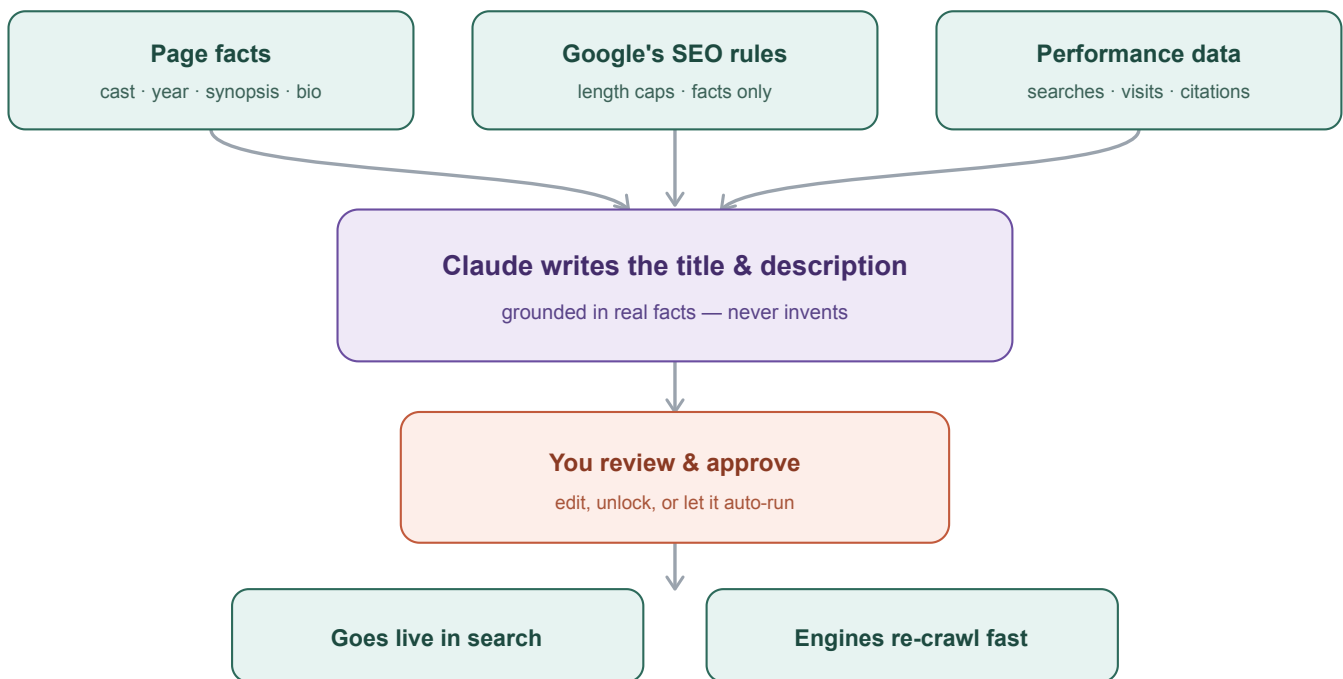
The rest of this handbook walks through each layer: **how a page gets optimized** (the writing), **where the site shows up** (the surfaces), **the AI tools** that do the thinking, and **the production toolkit** that runs the day-to-day.

- 1 — How a page gets optimized · the SEO Studio
- 2 — Where the site shows up · Google, AI, voice, image/video, social
- 3 — The intelligence layer · AI tools that think and write
- 4 — The production engine · keeping the catalog accurate at scale
- 5 — The flywheel · how it compounds over time

SECTION 1

How a page gets optimized

Every page needs two short pieces of text: the clickable headline and the one-line blurb you see under it in search results (and when a link is shared). There are about 6,800 pages on the site. The **SEO Studio** writes and maintains all of them.



Facts + Google's rules + your live performance data → the AI writes it → you review (or let it run) → it goes live and the search engines are pinged.

What's inside the SEO Studio

- **AI drafting, grounded in real facts.** For each page it gathers what the site already knows — a film's cast, year, synopsis; a person's bio and which film they're in; a journal post's actual text — and asks Claude (the AI) to write the best headline and blurb under strict, Google-based rules. It can only use real facts; it never invents a credit or an award.

- **Locks on the things that shouldn't change.** Your film titles follow a fixed format ("The Sound of Identity (2021) · Documentary by James Kicklighter"). The system is physically locked from rewriting those (and your library and director pages) — but you can click "unlock" to hand-edit any one.
- **Live performance data folded in.** It quietly pulls what people actually search to find each page (Search Console), how many visits it gets (Analytics), Bing's figures, and whether AI assistants cite it — and tunes the wording toward what's genuinely bringing people in. It refreshes that data on its own.
- **Suggest-only, with optional autopilot.** By default it shows you a suggestion to approve. For journal posts, photos, courses, press, and process documents you can let it write and publish automatically the moment you hit publish.
- **Never redoes finished work.** An "optimize only un-optimized" control skips pages that are already done, so completed work is left untouched.
- **A monthly check that only acts when it matters.** Once a month it looks at whether the search demand for a page has meaningfully shifted, and only then refreshes it. Stable pages are left alone.
- **Export / re-import.** Every page's title, description, and source facts can be exported to a spreadsheet for a bird's-eye audit, then re-imported.

SECTION 2

Where the site shows up

A great listing means nothing if it isn't published where people look. The site speaks the native language of **five different kinds of search** at once. Here's each surface and the custom machinery that feeds it.

Google & Bing — traditional search

- **Custom sitemap.** A complete, prioritized map of every page handed straight to Google and Bing — with freshness dates that update the moment a page changes, so updates are noticed in days, not weeks.
- **Structured data (schema).** Every page publishes a hidden, machine-readable "fact sheet" (a film's director, year, rating; a person's bio; an article's author and date) so Google can show rich results and understand exactly what each page is.
- **Smart titles & link previews.** Each page type gets a tailored, length-correct headline that won't get cut off in results — never generic boilerplate, never the brand name twice.
- **Snippet & preview controls.** Tells Google it's welcome to show large image previews and full text snippets — and quietly hides thin or duplicate pages so crawl attention concentrates on the real content.
- **Canonical addresses.** Each page declares its one true URL, so tracking links and dynamic pages never split into "duplicate content" that dilutes ranking.

- **RSS feeds.** Standard subscribe-able feeds for the journal and press, for readers and aggregators.

AI & generative search — ChatGPT, Perplexity, Claude, Gemini

This is the newest and least-common surface — most sites do none of it. The site publishes a whole set of files written specifically for AI engines to read.

- **lms.txt.** A clean summary index — your filmography, press, services, credentials, and official identifiers — so an AI can grab who you are and what's new without wrestling with the webpage.
- **lms-full.txt.** The complete corpus in plain text (every film, recent journal posts, FAQs) for the AI engines that prefer to ingest everything at once.
- **about.json.** A machine-clean biography — name, birthplace, job titles, pronunciation, official IDs, current status, representation — delivered in one request for AI pipelines that prefer pure data.
- **ai.txt & an AI-bot allow-list.** An explicit, friendly "yes, you may read and cite this" policy naming every major AI crawler (GPTBot, ClaudeBot, PerplexityBot, Google-Extended...) so they're permitted, not blocked.
- **AI plugin manifest.** A discovery file (in the .well-known folder) that tells AI agents exactly where to find your data feeds.
- **Named-author & freshness signals.** Clear authorship and "last updated" markers — research shows AI engines cite named, recently-updated content several times more often.

Voice search — Siri, Alexa, Google Assistant

- **Pronunciation built in.** The machine bio spells out how to say "Kicklighter," so voice assistants read the name correctly aloud.
- **Identity anchors.** Official Wikipedia / IMDb / Wikidata IDs are published so an assistant knows *which* James Kicklighter is being asked about.
- **FAQ-shaped answers.** Frequently-asked questions are published in a structured, answerable form — the exact shape voice assistants pull from.
- **Process notes & director's statements** are surfaced in the AI files, so "what was James's process for X?" has a real answer to read.

Image & video search — the visual catalog

- **Image sitemap.** A dedicated map of every behind-the-scenes and production photo, organized by film and event, so Google Image Search can find the whole catalog.
- **Video sitemap.** Every trailer (YouTube/Vimeo) is indexed with its thumbnail, length, and description for Google Video Search.

- **Photo-gallery fact sheets.** Each film's photo page publishes structured data tying every still — with caption, photographer, and filming location — to the film it documents.
- **Dedicated photo landing pages.** Each film gets a "{Film} — Photos" page tuned to surface large image previews in search.

Social sharing — Facebook, Twitter/X, LinkedIn, Slack, Discord

- **Dynamic share previews.** When any link is shared, the preview shows the real title, the real synopsis, and the right image — never a bare site name.
- **OG Card Studio.** A one-click generator for branded 1200×630 share cards in your brand fonts, customizable per page, downloadable or attached to a page in one action.

Speed, because speed is ranking

Google and AI engines both favor fast pages. The site is tuned aggressively under the hood: the critical stylesheet is preloaded; connections to image, video, and analytics services are warmed up only when needed; analytics scripts wait until a visitor actually engages (keeping the page instant); and unused code is stripped per page. Visitors get a fast site; search engines reward it.

SECTION 3

The intelligence layer

These are the tools that **think** — they use AI to watch how the site is performing, find what's missing, and do the heavy writing. This is where the site stops being a brochure and starts being a system that improves itself.

The AI Citation Tracker — are the AI assistants citing you?

A first-of-its-kind feedback loop: it asks the AI assistants questions about you and your films, records whether they cite you, and turns the answers into a to-do list.

- **Citation Probe.** On an automatic weekly schedule it asks ChatGPT, Claude, Perplexity, and Gemini a curated set of questions about you and your work, and records whether each one cites you.
- **Content-Gap Analysis.** It finds the questions AI engines clearly field about your field that your site doesn't yet answer — so you know exactly what to publish next.

- **Citation Suggestions.** It turns the data into concrete, one-click ideas: expertise topics to add, FAQs to answer, descriptor phrases the AIs are already using for you.
- **Competitor Detection.** It tracks which filmmakers get mentioned alongside you, separating intentional peers from automatically-spotted adjacent names.
- **Query Candidates & Clusters.** It mines the AIs' own "people also ask" follow-ups for new questions to track, and shows which topics and peers cluster together — useful for positioning.
- **Citation Trends.** Little trend-lines per question and per engine show whether your citation rate is climbing — so you can see if your changes are actually working.

AI that writes & enriches

- **SEO Studio** (Section 1) — the AI that writes every page's title and description.
- **Schema Enrichment.** The AI reads each film's synopsis and auto-fills the behind-the-scenes "fact sheet" fields — themes, keywords, intended audience — that help both Google and AI engines place the work.

AI for the photo library

- **Duplicate Detection.** AI vision spots near-identical frames from a burst and keeps the best one, archiving the rest — so the galleries stay sharp without manual culling.
- **Photo Captioning & Alt Text.** AI vision writes a caption, description, and alt text for every behind-the-scenes still (which is also what makes them findable in image search).
- **Engagement Scoring.** AI scores each photo for visual interest, so the strongest images automatically rise to the front of the On Location gallery.
- **Contact Sheet Planner.** A weekly social scheduler that proposes five photos, balances the mix, and AI-writes captions tailored for Instagram, Threads, LinkedIn, and Facebook.
- **"Best of James."** Surfaces every photo where you've been identified, ranked by quality — an instant shortlist for press kits and profiles.

Notifying the world automatically

- **IndexNow.** The instant any page's wording changes, the site pings Bing and Yandex to re-crawl it within minutes — and stamps the sitemap so Google notices too. No manual submitting.

SECTION 4

The production engine

Behind the public site is a purpose-built production system that keeps a large, fast-growing catalog accurate and richly documented — thousands of photographs, a deep web of people and

credits, and years of press — and keeps it all clean automatically.

An industrial photo operation

The behind-the-scenes archive runs to thousands of images. It stays curated end to end: near-identical frames from a burst are resolved down to the best one, every photo is described for search and accessibility, and the strongest images are surfaced first — so the galleries stay sharp without hand-culling, and the entire library is findable in image search.

A living credits graph

Every collaborator, cast member, and creative influence is connected into one deduplicated web, with verified biographies and official identifiers attached. The relationships between people and projects are made explicit and machine-readable — which is what lets search engines and AI assistants understand exactly who did what.

Press, attributed automatically

Years of coverage are organized into a live, structured library — each article matched to the work it covers and attributed to its publication. Credibility signals are captured consistently and kept current, rather than assembled by hand.


Accurate, fast publishing

New films and posts begin from authoritative sources with their core facts pre-filled, and the catalog is continuously checked for gaps and inconsistencies. The result is a site that grows quickly while staying correct, complete, and consistent across every page.

SECTION 5

The flywheel

The reason it all matters is that these pieces aren't separate — they feed each other in a loop that gets stronger every time you publish.



...then the AI Citation Tracker measures what landed, finds the gaps, and feeds the next thing to publish.

Publish → the AI optimizes → the engines are notified → you get found → the tracker measures and suggests what's next. Every cycle compounds.

Most websites are static — you build a page and hope. This one is a **living system**:

- **You publish** a film, a journal entry, a batch of photos.
- **It optimizes itself** — the AI writes the search listing from the real facts and your live traffic data.
- **It announces itself** — the search engines are pinged to re-crawl within minutes.
- **It gets found** — across Google, AI assistants, voice, image and video search, and social.
- **It measures itself** — the citation tracker checks whether the AI engines actually picked it up, and the analytics show real traffic.
- **It tells you what's next** — gaps and suggestions become your next thing to publish, and the loop runs again.

The net effect: the work you've already done keeps working — quietly, automatically, and across every surface where people and machines go looking — so the films, the writing, and the photography reach the largest possible audience with the least possible upkeep.