

The SEO Francisco AI Search And SEO Guide For Claude

15 skills, scheduled audits, prompt packs, and release-day QA gates for evidence-first
SEO work with Claude.

Version 2026-07-03

SEO Francisco

Start Here

Claude can make SEO work faster, but only when the workflow is grounded in evidence. A polished report is not enough. Before you trust the answer, ask:

- What data did Claude see?
- What rule decided the recommendation?

If the workflow cannot answer those questions, the output is not ready for a client, roadmap, or publication decision.

The Operating System

Use four layers:

- Evidence first.
- Skills second.
- Schedule third.
- QA last.

Evidence first means Claude receives source packs, crawl exports, Search Console data, GA4 data, logs, AI answer snapshots, cited domains, raw transcripts, and notes from real review.

Skills second means repeatable jobs have stable inputs, rules, and outputs.

Schedule third means the checks happen daily, weekly, monthly, and on release day.

QA last means the final output is checked for missing sources, invented metrics, weak confidence labels, AI-tells, internal links, image quality, video quality, and publication parity.

The 15 Skills

1. Brand Brief Keeper

Keeps the core brand context Claude should read before every SEO task.

Inputs:

- Priority services and money pages

- Target customers
- Tone rules
- Competitors
- Claims that need evidence
- Claims to avoid

Output:

- A one-page working brief
- Do-not-say list
- Source requirements

2. Source Pack Builder

Turns research into a clean evidence bundle.

Inputs:

- Official docs
- Research papers
- product docs
- crawl data
- first-party examples

Output:

- Source summary
- Confirmed facts
- Unproven claims
- Links for citation
- Draft boundaries

3. AI Citation Gap Finder

Finds prompts where competitors appear in AI answers and your brand does not.

Inputs:

- Prompt set
- Raw answers from ChatGPT, Claude, Perplexity, Gemini, AI Mode, or another tracked surface
- Brand and competitor list

Output:

- Missed prompts
- Cited competitor domains
- Probable content gap
- Next page or update to build

4. Competitor Citation Teardown

Breaks down the source types AI systems cite for money prompts.

Inputs:

- Raw AI answers
- Cited URLs
- Competitor mentions

Output:

- Domains cited most often
- Page format patterns
- Evidence style
- Content gaps
- Link/distribution targets

5. Prompt Fan-Out Mapper

Expands a single commercial prompt into the supporting questions an AI answer needs.

Inputs:

- Money prompt
- Target page
- Existing headings
- Competitor pages

Output:

- Main prompt
- Support prompts
- Missing subtopics
- Headings to add or remove

6. Content Decay Checker

Finds pages where performance or freshness is slipping.

Inputs:

- GSC clicks and impressions
- Query data
- Last updated date
- SERP notes

Output:

- Pages down by threshold
- Lost query groups
- Refresh priority
- Evidence to add

7. Internal Link Opportunity Mapper

Finds relevant internal links using page topic, entity, and user journey.

Inputs:

- New draft
- Existing page list
- Target service pages
- Related articles

Output:

- Anchor suggestions
- Source page
- Target page
- Reason for link

8. Schema Evidence Auditor

Checks whether structured data matches visible page facts.

Inputs:

- Page HTML
- rendered DOM
- schema JSON-LD
- Rich Results Test output when available

Output:

- Confirmed schema issues
- Unsupported markup
- Missing schema opportunities
- Validation steps

9. Rendering And JavaScript SEO Inspector

Compares source HTML, rendered HTML, and crawler-visible content.

Inputs:

- raw HTML
- rendered DOM
- screenshots
- crawl output

Output:

- content present in source
- content present after render
- missing links
- blocked resources
- test plan

10. AI Answer Snapshot Logger

Saves raw AI answer evidence before interpretation.

Inputs:

- platform
- model or surface
- prompt
- answer
- timestamp
- citations

Output:

- raw answer archive
- mention counts
- cited domains

- caveats

11. Entity And Author Evidence Gate

Checks whether the page shows why this author or brand should be trusted.

Inputs:

- author bio
- source evidence
- examples
- credentials
- related pages

Output:

- missing experience signal
- missing citation
- weak author context
- improvement notes

12. YouTube SEO And Video Citation Mapper

Connects article, video, title, description, chapters, and related videos.

Inputs:

- final video
- YouTube title
- description
- transcript
- article URL
- related video URLs

Output:

- title options
- description with article and video interlinks
- chapters
- pinned comment
- embed notes

13. Release-Day QA Gate

Checks whether a page is actually ready to publish.

Inputs:

- EN post
- ES post
- images
- video
- links
- build output
- live URL checks

Output:

- PASS/FAIL
- blockers
- exact files or URLs to fix

14. Weekly AI Visibility Report

Turns raw tracking into three decisions.

Inputs:

- prompt snapshots
- citation changes
- GSC movement
- competitor movement
- content releases

Output:

- visibility summary
- top movers
- citation gaps
- three priorities

15. Memory And Decision Log Keeper

Prevents the workflow from forgetting corrections.

Inputs:

- user corrections

- failed attempts
- QA blockers
- approved assets
- final URLs

Output:

- memory note
- package ledger update
- next-action summary

Prompt Shape

Every prompt should follow the same pattern:

```
Role:  
Context:  
Data:  
Rules:  
Output:
```

The rules matter most.

Use rules like:

- Use only the provided evidence.
- If data is missing, say "data needed."
- Do not invent traffic, rankings, citations, revenue, or dates.
- Label confidence as confirmed, likely, or needs validation.
- Separate facts from recommendations.
- Name the source behind each claim.
- Prioritize by business impact and search risk.

Daily Checks

Run these every working day:

- Check blocked publication tasks.
- Check pages with sudden GSC drops.
- Check indexing or crawl blockers.
- Check new AI answer snapshots for priority prompts.
- Check whether any generated draft needs humanization before publishing.

Weekly Checks

Run these once a week:

- AI citation gap audit.
- Competitor citation teardown.
- Content decay review.
- Internal link opportunity review.
- Schema and rendering spot checks.
- Weekly report with three priorities.

Monthly Checks

Run these once a month:

- Refresh the brand brief.
- Refresh prompt sets.
- Review source drift.
- Review schema quality.
- Review crawl patterns and JavaScript rendering.
- Review content clusters for scaled sameness.

Release-Day Gate

Before publishing a page:

- Source pack exists.
- Draft uses confirmed sources.
- AI-tells audit is clean or humanized.
- Internal links are useful.
- Related articles are added.
- Images add concrete value.
- Feature image and infographic pass QA.
- Video script passes preflight.
- VoxCPM2 narration passes audio QA and listening.
- HeyGen raw is Avatar III and 1080p.

- Final video is edited and uploaded.
- YouTube description links article and related videos.
- AI use disclosure is set to No unless explicitly changed.
- EN and ES pages publish together.
- Build passes.
- Live URLs return 200.
- IndexNow is submitted.
- Search Console indexing is requested when available.
- Memory is updated.

The Five-Skill Starter Set

If you want to start small, use only these five:

- Brand Brief Keeper
- AI Citation Gap Finder
- Content Decay Checker
- Internal Link Opportunity Mapper
- Release-Day QA Gate

Those five prevent most AI-assisted SEO workflows from becoming polished guesswork.

Final Rule

Claude is strongest when it has real context, a narrow job, and permission to say "I do not have the data." Build the workflow so that answer is allowed.