Fusionn

THE 3 KEYWORD SKILLS · 2026

Find your keywords with Claude

The three skills I use with Claude to pull out my keywords, group them into pages and isolate the ones that bring in a client, ready to copy straight into your skills folder.

I've already shared the skills that cover my whole SEO workflow. For keywords, there are three specific skills: the first runs a general analysis around the main keyword, the second groups the keywords by intent so you know which page to create, and then the third ranks them by business potential. The idea is to create ultra-high-quality pages that are useful to the user, with the goal of capturing a qualified email address to engage them with your brand.

Here are the three skills to copy into Claude. Each one is documented the same way: what it does, when to trigger it, and the full SKILL.md to paste into your .claude/skills/ folder.

A keyword skill can replace a tool

Producing a keyword list has become a commodity. Any LLM or SEO tool will give you a complete list in a few seconds. What matters is therefore finding keywords that won't be eaten by AI: queries that let you rank on business keywords and, if possible, less competitive ones.

As you'll have gathered, the type of keyword we go after has changed too. AI answers informational queries for you, so we aim for decision-stage queries (action + transactional), meaning every query where the user wants to do something, not just read a piece of content (for that, they'll go to ChatGPT). That doesn't mean no informational content at all, because obviously you still need to add content for your semantics, your internal linking, to share your values, and so on.

HERE'S HOW I ACTUALLY USE IT

1. I give it as much context as I can: my Google Search Console exports, my client's tone of voice, sales call scripts, support tickets, and so on. Without that context, Claude guesses, and it guesses the same thing it would for everyone else. So it's hard to be original.
2. I ask it for the queries a user types when making a decision, not the informational ones that AI will eat. If someone just wants to read, ChatGPT does the job and I don't go there. If they want to compare, calculate or choose, that's where I create the page.
3. I force it to sort by business intent, never by volume. Most of the keywords that bring in a client have no volume at all: they're micro-intents that traditional SEO tools struggle to find.
4. I cross-reference it with my proprietary data to do better than the average on the SERP. Data that exists elsewhere is a commodity; data that exists only with me is what sets me apart and lets me rank.

BEFORE RUNNING THE FIRST SKILL

The context you give Claude

Without real client context, your vector stays generic. Same vector as everyone else = nowhere in the SERP. Context makes the vector diverge, and the divergence creates the Surprise Gap.

Four sources feed the context I inject into Claude:

CLIENT CALLS

Verbatim quotes, real objections, prospects' vocabulary. Everything that surfaces neither in the SERP nor in SEO tools. Source number one: client data, which becomes proprietary data that nobody else can have. Every sales call contains 10 to 20 exact phrasings your prospects use to talk about their problem. Those are our new keywords.

MARKET RESEARCH

Your industry's figures: trends, market shares, market size. That's what lets you put forward verifiable data instead of mere opinions. A page that cites real figures is picked up more easily by Google and AI than a page that just gives an opinion.

REDDIT

Users' real questions, their concrete problems, their experience reports. On Reddit, users talk to other users: so you need to capture what is said about your product or service on this type of forum, then use it on pages to address micro-intents that a classic approach can't see.

GROK OR PERPLEXITY

Fresh data from the users themselves, social signals. The complement an average piece of content doesn't have, because it stops at queries everyone already covers. It also helps you find emerging topics in your field, so very handy.

One ranking factor on LLMs is the Surprise Gap: the measured gap between your page and the corpus average for the query.

Perplexity also lets you run in-depth, qualitative market research to feed an editorial angle, a query or your content.

Skill #1 · Keyword research

seo-recherche-mots-cles

DEFINITION

Start from a topic or a few seed keywords and produce an exhaustive list, qualified by intent and prioritised. It's the first step: until this list exists, clustering and briefs stay too generalist and no longer let you rank on competitive keywords.

WHY YOU NEED IT

This is where your context / proprietary data becomes extremely important (the one you gathered in the previous step). Without it, you just repeat what's already said on the SERP. Like everyone else. The skills I give you let Claude produce user-based queries, drawn from your clients' real vocabulary.

And one important rule in each skill: Claude isn't allowed to invent or hallucinate (you still need to check, obviously). It will structure the approach and handle search intents.

THE FULL SKILL

```
---
name: seo-recherche-mots-cles
description: |
  Keyword research from scratch: start from a topic or a few seed keywords and produce an exhaustive,
qualified and prioritised list — each keyword with its intent, volume and difficulty. 5-step pipeline:
frame the topic → semantic expansion → qualification (intent/volume/difficulty) → noise filtering →
prioritised table. Strict anti-hallucination rule: no volume or difficulty is ever invented, the [TO
SOURCE] placeholder is mandatory.

  ALWAYS use this skill when the user says: "keyword research", "find me keywords", "list of keywords
on [topic]", "which keywords to target", "explore the queries on", "I have a topic, give me the
keywords", "keyword expansion".

  First building block of the keyword workflow: its output feeds seo-clustering-mots-cles, then seo-
mots-cles-decisionnels.
---

# Skill — Keyword Research

## When to trigger

Start from a topic or seed keywords and produce an exhaustive, qualified and prioritised keyword list.
It sits upstream of everything: until the list exists, neither clustering nor the brief has anything
to work with.

## Doctrine

- AI **never** invents a search volume or a difficulty score. It structures, qualifies intent,
suggests variants. The figures come from a tool: GSC, Ahrefs, Semrush, Keyword Planner.
- We don't start from competitors. We start from the audience's intent and from the offer.
- A keyword without a clear intent isn't a keyword, it's noise.
- Exhaustiveness first, sorting second: generate broadly, filter afterwards.

## Required input

| Source | Mandatory |
|--------|-----------|
```

```
| Topic or 3-5 seed keywords | Yes |
| Sector + audience / persona | Yes |
| Offer (what we sell, conversion point) | Yes — otherwise the funnel can't be qualified |
| Geographic market + language | Yes |
| GSC / Ahrefs / Semrush export if available | Recommended — makes the volumes real |

## Pipeline (5 steps)

1. **Frame** — topic, perimeter, what is in scope and what is out.
2. **Semantic expansion** — for each seed, generate:
   - modifiers (how, why, price, reviews, best, vs, without…)
   - questions (People Also Ask patterns)
   - long tail (3 words or more)
   - co-occurrences and related entities
   - synonyms and lexical variants
   Goal: **at least 50 to 150 raw keywords** before sorting.
3. **Qualify each keyword** — intent (Know-Simple / Know / Do, see `seo-cluster-aeo`), volume (`[TO SOURCE]` or the real figure if an export is provided), estimated difficulty (proxy: number of words, presence of brands, SERP type), funnel stage.
4. **Filter the noise** — remove off-topic keywords, those off-intent with regard to the offer, strict duplicates, irrelevant competitor branded queries.
5. **Deliverable** — prioritised table.

## Mandatory output

```
KEYWORD RESEARCH — '[Topic]'

| Keyword | Intent | Volume | Difficulty | Funnel | Note |
|---|---|---|---|---|---|
| ... | Know | [TO SOURCE] | Medium | Know | ... |
| ... | Do | 1,300 (Ahrefs) | High | Decision | ... |

Summary: [total] keywords · [n] Do · [n] Know · [n] Know-Simple
Top 5 to prioritise: [list]
```

Absolute rules

- Zero invented volume or difficulty → `[TO SOURCE]`.
- If a tool export is provided, use **its** figures: never round them or estimate them at random.
- Generate broadly (50+ before filtering), deliver sorted.
- Every keyword has an identified intent, or it drops off the list.

Workflow sequence

`seo-recherche-mots-cles` (this list) → `seo-clustering-mots-cles` (groups into pages) → `seo-mots-cles-decisionnels` (isolates the ones that convert) → `seo-brief-contenu` / `seo-cluster-aeo` / `seo-programmatique-pseo`.

Related concepts

`know-simple-know-do` · `funnel` · `rrf` · `clustering-mots-cles` · `mots-cles-decisionnels` · `brief-contenu`
```

Skill #2 · Keyword clustering

seo-clustering-mots-cles

DEFINITION

Turn a raw list of N keywords into usable clusters, where each cluster becomes a page.

WHY YOU NEED IT

A flat list of 150 keywords is unusable. The real criterion for grouping them isn't how similar the words are, it's shared SERPs: if Google shows the same top 10 for two queries, it treats them as the same intent, so one page.

Get that wrong and you create two pages on the same intent and set up your own cannibalisation. It's the first cause of stagnation I see with my clients. The rule fits on one line: one cluster, one page, one pivot keyword.

THE FULL SKILL

```
---
name: seo-clustering-mots-cles
description: |
  Groups a raw keyword list into usable clusters — 1 cluster = 1 page. Grouping criterion: shared
SERPs (two keywords that display the same top 10 = the same intent = a single page). 5-step pipeline:
import the list → normalise (deduplicate, clean) → group by SERP intent → name each cluster and pick
the pivot keyword → assign 1 cluster = 1 page with cannibalisation detection.

  ALWAYS use this skill when the user says: "cluster my keywords", "group this keyword list", "keyword
clustering", "group my queries", "I have 200 keywords to organise", "which keywords go on the same
page", "turn my list into pages".

  Differs from seo-cluster-aeo, which starts from ONE pillar keyword to build an architecture. Here we
start from a RAW LIST of N keywords and cut it into pages. Architecture building block of the keyword
workflow: after seo-recherche-mots-cles, before seo-brief-contenu.
---

# Skill — Keyword Clustering

## When to trigger

Turn a raw list of N keywords into usable clusters, where each cluster becomes a page.

## Doctrine

- The real SEO clustering criterion = **shared SERPs**. If Google displays the same top 10 for two
queries, it treats them as the same intent → a single page.
- Semantic proximity alone isn't enough. "Cheap car insurance" and "car insurance rates" look alike
but can have distinct SERPs.
- 1 cluster = 1 page = 1 pivot keyword. Two pages on the same cluster = cannibalisation set up in
advance.
- Without access to the real SERPs, AI clusters by intent + semantics **as a proxy**: the output is a
hypothesis to validate, never a truth.

## Required input

| Source | Mandatory |
|--------|-----------|
| The raw keyword list (paste, CSV, output of `seo-recherche-mots-cles`) | Yes |
| Intent + volume per keyword | Recommended |
```

```
| SERP overlap data from a tool (Keyword Insights, Ahrefs "Parent topic") | Recommended — makes the clustering reliable |
| The site / the offer (to guide naming and detect cannibalisation) | Recommended |

## Pipeline (5 steps)

1. **Import the list** — accept a raw paste, a CSV or the output of `seo-recherche-mots-cles`.
2. **Normalise** — deduplicate, remove the noise (keywords without intent, out-of-scope competitor branded queries), merge strict spelling variants (singular/plural, accents).
3. **Group** — criterion #1: SERP overlap if the data is provided. Otherwise proxy: identical intent + same semantic head. Each group = one cluster.
4. **Name + pivot** — for each cluster: name (head term), dominant intent (Know-Simple / Know / Do), pivot keyword (the most representative one, with the highest volume) + secondary keywords.
5. **Assign** — 1 cluster = 1 page. Detect the clusters that overlap an existing page → cannibalisation risk, refer to `seo-cannibalisation`.

## Mandatory output

```
CLUSTERS — '[List / topic]'

| Cluster | Pivot keyword | Secondary | Intent | Page (new / existing) | Cannibalisation risk |
|---|---|---|---|---|---|
| ... | ... | ... | Know | New | — |

Not clustered: [orphan keywords left on their own]
To validate: confirm each cluster with a real SERP overlap (tool or manual check).
```

Absolute rules

- 1 cluster = 1 page = 1 pivot. Never two clusters on the same intent (MECE principle).
- A keyword belongs to a single cluster only.
- Without SERP data, explicitly mark the output "hypothesis — to be validated by real overlap".
- Don't force an orphan keyword into a cluster: leave it as "not clustered".
- If a cluster overlaps a live page → flag cannibalisation, don't create a duplicate.

Workflow sequence

After `seo-recherche-mots-cles` → each cluster becomes an input: `seo-brief-contenu` (1 cluster = 1 brief), or `seo-cluster-aeo` if the cluster is big enough to become a pillar. Do-intent clusters → `seo-mots-cles-decisionnels`.

Related concepts

`serp-overlap` · `mece` · `cannibalisation` · `cluster-aeo` · `passage-ranking` · `maillage-interne`
```

Skill #3 · Decision-stage keywords

seo-mots-cles-decisionnels

DEFINITION

Isolate, within your list, the queries that bring in a customer, not just traffic.

WHY YOU NEED IT

AI is eating informational queries. ChatGPT and the AI Overview answer "what is X" without a click to your site. The keywords that survive are the ones where the user wants to do something: compare, calculate, request a quote, choose. That's what I call action keywords, both decision-stage and transactional. They're the only ones that generate revenue in SEO today.

Before going after a keyword, put it through the two-question test: can ChatGPT answer it, and can it do better than you? If the answer to both is yes, the page is dead before it exists. Decision-stage keywords are spotted by their modifiers (price, quote, best, reviews, alternative), not by their volume.

THE FULL SKILL

```
---
name: seo-mots-cles-decisionnels
description: |
  Isolates decision-stage keywords — the transactional, bottom-of-funnel queries that convert, as
opposed to informational queries that only bring traffic. 5-step pipeline: define the offer and the
conversion point → spot the decision signals → classify by decision stage → score the conversion
potential → prioritised shortlist with recommended page and format. Starts from an existing list or
from a topic.

  ALWAYS use this skill when the user says: "decision-stage keywords", "transactional keywords",
"keywords that convert", "bottom-of-funnel queries", "Do queries", "bottom funnel keywords", "which
keywords make money", "sort my list by conversion potential", "commercial keywords".

  Conversion building block of the keyword workflow: plugs in after seo-recherche-mots-cles or seo-
clustering-mots-cles.
---

# Skill — Decision-Stage Keywords

## When to trigger

Isolate, within a list or a topic, the queries that bring in a customer — not just traffic.

## Doctrine

- Traffic isn't the goal, conversion is. A keyword with 10,000 views a month and no buying intent is
worth less than a keyword with 80 views that brings in a customer.
- A decision-stage keyword = the internet user is in the middle of choosing, comparing or buying. They
have mentally taken out their credit card.
- Decision-stage keywords are spotted by their **modifiers**, not by their volume.
- Bottom of funnel ≠ necessarily low volume: "best [category]" can carry a lot of weight.
- Decision-stage keywords are handled first: it's the shortest ROI in SEO.

## Required input

| Source | Mandatory |
|--------|-----------|
```

Keyword list OR topic	Yes
Offer: what we sell, price, conversion point (form, purchase, demo, call)	Yes
Audience / persona	Yes
Competitors (for branded queries)	Recommended

Decision signals (modifier table)

Category	Modifiers
Direct purchase	buy, order, price, rates, quote, subscription, sign up
Comparison	best, top, vs, comparison, alternative to, which to choose
Evaluation	reviews, test, feedback, reliable, scam
Local / immediate	near me, [city], online, delivery, open
Brand	[your brand], [your brand] reviews/price, [competitor] alternative

Decision stages

- **Do — Purchase**: the internet user wants to convert now.
- **Do — Comparison**: they're choosing between options, you included.
- **Do — Local**: buying intent + geographic constraint.
- **Do — Brand**: they already know you, or know a competitor.
- **Commercial Know**: "how to choose an X" — pre-decision, to capture early.

Pipeline (5 steps)

1. **Define the offer** — what we sell, the price, the real conversion point.
2. **Spot the signals** — run the list or the topic through the decision modifiers filter.
3. **Classify by decision stage** — sort each retained keyword.
4. **Score the conversion potential** — proximity to the offer (1–5) × buying intent (1–5) × feasibility of ranking (1–5). Flag competitor branded queries to check their legitimacy.
5. **Deliverable** — prioritised shortlist; for each keyword: target page (landing page, comparison, product page, local page), format, CTA.

Mandatory output

```
DECISION-STAGE KEYWORDS — '[Offer]'

| Keyword | Stage | Conv. score | Target page | Format | CTA |
|---|---|---|---|---|---|
| ... | Do — Comparison | 64/125 | Comparison page | Table + verdict | Quote |

Top 3 to go after first: [list + justification]
```

Absolute rules

- Never classify a purely informational keyword as decision-stage to pad the list.
- Competitor branded query → check legitimacy: honest comparison page, no deception.
- Every decision-stage keyword points to a page with an offer + CTA, otherwise it doesn't convert.
- Prioritise proximity to the offer + feasibility, not raw volume.

Workflow sequence

After `seo-recherche-mots-cles` or `seo-clustering-mots-cles` → feeds `seo-product-led-seo` (tools for Do queries) and `seo-brief-contenu` (conversion pages).

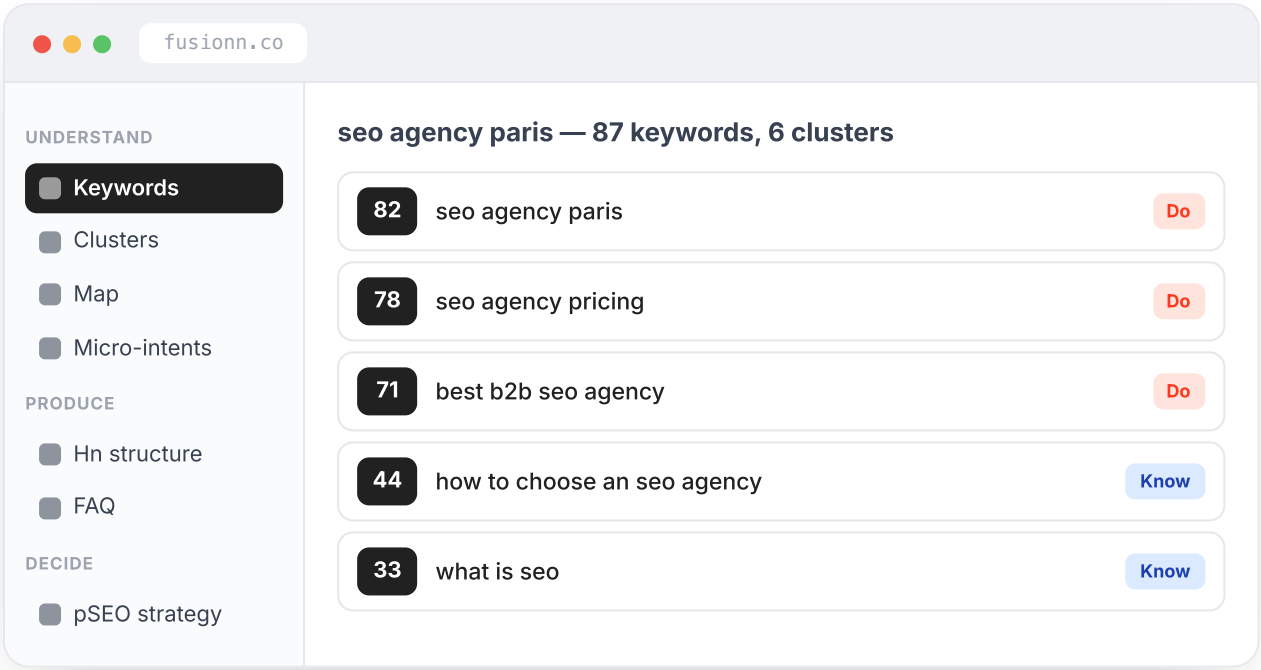
Related concepts

`know-simple-know-do` · `funnel` · `product-led-seo` · `fully-meets` · `cannibalisation` · `e-e-a-t`

GOING FURTHER

For those who want to push the analysis across 14 angles

Running the three skills by hand in Claude is enough to get your keywords. Fusionn goes further: starting from a topic, it breaks your subject down across 14 angles, from clusters to micro-intents, from the Hn structure to objections, all the way to the pSEO strategy. Each keyword is scored on its business potential and sorted by intent.



Interface preview. Each keyword is scored on its business potential and sorted by intent. The same logic as the three skills, but with nothing to copy and paste.

Try it on one of your keywords

You give a topic, you get the qualified list, the clusters and the shortlist of what converts. Three free analyses, no credit card.

Open Fusionn



Timothée Boussardon

SEO consultant for 6 years. I help B2B companies turn their SEO into an acquisition channel, with an SEO-AI system plugged into their real data (Search Console, sales calls, support tickets).

Top 3 · golf ball

Top 3 · Qualiopi consulting

Top 3 · aparthotel Bordeaux

Follow me on LinkedIn →