


The *Findable* Business

Plain-English SEO for Google and
AI Search



The Findable Business

Plain-English SEO for Google and AI Search

DAVN00T

The Findable Business: Plain-English SEO for Google and AI Search

© 2026 Davnoot. All rights reserved.

No part of this book may be reproduced or transmitted in any form or by any means without written permission from the publisher, except for brief quotations in reviews and articles, and except for the checklist in Appendix A, which you may print and share freely with credit to davnoot.com/seo-checklist.

Published by Davnoot, Montreal, Quebec, Canada. davnoot.com

First edition, 2026.

Disclaimer. Search engines and AI assistants change how they work all the time, often without notice. This book describes how things stood when it was written, based on public documentation and our own experience. Nothing in it guarantees any ranking, citation, traffic or revenue, and no honest person can promise you those. Where the book touches on law, including Quebec's Charter of the French language and Law 25, it gives general context only. It is not legal advice; for a formal opinion, talk to a lawyer.

The businesses used as examples in this book are fictional composites. Any resemblance to a real business is a coincidence. Product names and trademarks belong to their owners.

How to use this book

This book is for people who run or market a business and need it to be found. You might own a clinic, a trades company, a law practice, a shop, or a software product. You probably have a website that someone built for you, a Google Business Profile you set up once, and a nagging sense that you should be doing more with both.

You do not need to become an SEO specialist to get most of the value from search. For a small or mid-sized business, most of what moves results is unglamorous and well understood. Pick the right searches. Give each one a page that deserves to win, make sure search engines can read that page, and earn some trust from the rest of the web. AI search added a few new habits to that list. It did not replace the list.

The book follows the order we use with our own clients. The early chapters are about decisions: which searches to chase, and which page should own each one. The middle chapters cover work on the pages and on the site itself. After that come local search, links, AI assistants, structured data, bilingual sites and measurement. The last chapter turns all of it into a 90-day plan.

Before you start, know three things.

- **The checklist is the spine.** Appendix A reproduces the Davnoot SEO Checklist, 44 checks in eight groups. When a chapter refers to "Checklist #12," it is explaining that item. You can work through the

checklist on its own, but the chapters tell you why each item matters and how to do it properly. A printable copy is free at davnoot.com/seo-checklist.

- **Every chapter ends with "Do this week."** These are three to five small actions, each tied to checklist items. Read only the chapter titles, do the weekly lists, and you will still be ahead of most of your competitors.
- **Some of it will go out of date.** Search changes. When we describe a specific rule from Google or another platform, we say what it was when we wrote this. Check the source before you bet a quarter's budget on it.

We are a digital marketing agency in Montreal. We learned much of this book the hard way, on our own website, which runs in English and French. Where we made a mistake, we say so, and you are welcome to skip it.

Terms you may not know are defined in the glossary in Appendix B.

Contents

How to use this book	4
CHAPTER 1 How search works now	7
CHAPTER 2 Picking keywords you can actually win	13
CHAPTER 3 One page, one job	21
CHAPTER 4 On-page SEO that reads like a human wrote it	28
CHAPTER 5 Technical SEO without the panic	36
CHAPTER 6 Content that earns its ranking	46
CHAPTER 7 Local SEO	54
CHAPTER 8 Earning links honestly	62
CHAPTER 9 Getting cited by AI	71
CHAPTER 10 Structured data	81
CHAPTER 11 Bilingual SEO	88
CHAPTER 12 Measuring what matters	96
CHAPTER 13 Your first 90 days	103
APPENDIX A The Davnoot SEO Checklist	111
APPENDIX B Glossary	116
About Davnoot	120

How search works now

A homeowner in Longueuil wakes up to a cold house in January. The furnace is making a clicking sound and not producing heat. She does two things in the next ten minutes. First, she asks her phone why a furnace would click and not light. She gets a tidy answer about igniters and flame sensors, assembled by an AI from a handful of websites, and she does not click any of them. Then she searches "furnace repair Longueuil" and calls the second company on the map because it has more recent reviews and says it does emergency calls.

Most of what changed in search over the last few years is in that story. So is most of what didn't.

Three jobs every search engine does

Every search system, old or new, does the same three jobs.

1. **Crawling.** A program called a crawler (Googlebot, Bingbot, OAI-SearchBot and others) visits pages, follows links and downloads what it finds.
2. **Indexing.** The engine reads each page, works out what it is about, and stores a version of it in a huge database called the index. A page that is not in the index cannot appear anywhere, no matter how good it is.

3. Ranking and answering. When someone searches, the engine picks the most useful pages from its index and shows them. Traditional results show a list of links. AI features read several of those pages and write an answer, usually with links to the sources.

Google describes what it weighs when ranking in plain terms: the meaning of the query, the relevance of the page, the quality of the content, the usability of the page, and the context of the searcher (where they are, what language they use, what device they are on). None of those are secret, and each chapter in this book comes back to one of them.

The order matters. A page that cannot be crawled cannot be indexed, and a page that is not indexed cannot rank or be quoted. Plenty of business websites fail at step one or two and spend their money on step three. Chapter 5 is about making sure you are not one of them.

What changed

Google now answers some questions itself

Google's AI Overviews put a written summary at the top of many results pages, with links to the pages it drew from. Google's AI Mode goes further and turns search into a conversation. For simple questions ("what does a flame sensor do", "how long does teeth whitening last"), the searcher may get what they need without visiting any website.

That has hurt businesses that relied on informational blog traffic. If your plan was to attract thousands of visitors with general how-to articles and hope a few became customers, the math is worse than it used to be.

People search inside AI assistants

ChatGPT, Perplexity, Microsoft Copilot, Google Gemini and Claude can all search the web and cite sources. A growing share of people, especially for research-heavy purchases, start with one of them: "What should I look for in a bookkeeper for a small restaurant in Quebec?" or "Which CRM works for a ten-person agency?"

The assistants do not have a separate, magic web. When they need current information, they search. ChatGPT has its own crawler for search (OAI-SearchBot) and has also drawn on third-party search providers, Bing among them. Perplexity runs its own crawler. Copilot sits on Bing. Google's AI features sit on Google's index. So the question "how do I show up in AI answers?" starts with the same question as always: are you in the index, and are you the best source on the topic?

Being named counts now, even without a link

In a list of ten blue links, the only thing that mattered was your position. In an AI answer, being named can help a business even when nobody clicks. You want to be the company that follows "For emergency furnace repair on the South Shore, people often mention..." So what other sites say about you (directories, review sites, local news, industry associations, Reddit threads) matters more than it did, because AI answers are built from those sources too. Chapter 9 covers this in detail.

What didn't change

The searches that make you money still end in a click or a call

Look at the homeowner again. AI answered her informational question. Her commercial search, the one with a location and the word "repair" in it, still ended on the map, in a review, and on the phone.

In our experience, AI summaries touch the searches closest to a purchase the least. Someone who wants a dentist in Laval, a commercial lease lawyer in Montreal or a Shopify developer still has to choose a business and contact it. An AI can shortlist. It cannot fill a cavity. That is why this book spends most of its time on commercial pages, local search and trust, and little on blogging for traffic.

Google still runs on the same principles

Google has said plainly that the same SEO fundamentals apply to its AI features, and that you do not need to add any special markup or file to appear in AI Overviews. The pages it uses as sources are crawlable, indexed, useful and trusted. When an agency tells you AI search needs a completely new playbook, ask which part of the old one stopped mattering. The honest answer is usually "none of it, but some new habits sit on top."

Trust still has to be earned

Search engines have always tried to measure whether other people vouch for you: links from other sites, reviews, mentions, consistent business details across the web. AI assistants rely on the same kinds of signals. Writing "Montreal's most trusted plumber" on your own homepage counts for very little. Getting other people to say it elsewhere is what counts.

Two kinds of visibility

There are two outcomes to work toward, and you will measure them differently.

Ranked. Your page appears in Google's or Bing's results, or in the map pack, for a search. You can measure this directly in Google Search Console and Bing Webmaster Tools: impressions, clicks, average position, per query and per page.

Cited or mentioned. Your business or page is named or linked in an AI answer. This is harder to measure. You can check by asking the assistants yourself, by watching referral traffic from chatgpt.com, perplexity.ai and similar sources in your analytics, and by noticing what new leads say when you ask how they found you.

For most small businesses, ranked visibility still produces most of the leads, and the work that earns it is also most of the work that earns citations. So you do not need two separate strategies. You need one

strategy and two ways of checking it.

Bing matters again

For years, many businesses ignored Bing because its share of searches was small. That is now a mistake. Bing's index feeds Microsoft Copilot and has fed ChatGPT's search results, and other tools license it too. Bing Webmaster Tools is free, takes ten minutes to set up (you can import your site straight from Google Search Console), and shows you crawl and indexing problems Bing has with your site. It is Checklist #3 for that reason.

Three myths to drop now

"SEO is dead." People have announced this every year for twenty years. What dies, repeatedly, is a specific shortcut: keyword stuffing, link farms, thin city pages, mass-produced AI articles. The fundamentals outlive every one of them.

"You need special schema or an llms.txt file to show up in AI." Google says there is no special structured data required for its AI features. An llms.txt file is a proposed convention that some tools read and Google has said it does not use. Neither one hurts, and Chapter 10 explains what structured data is good for. Neither one, on its own, moves AI visibility.

"With AI answers, websites don't matter." AI answers are summaries of websites. If yours is not one of the sources, you are relying on other people to describe your business accurately. Your site is the one place on the web where you control the facts.

Where to start

Before you change anything, set up the free tools that show you what search engines see. Google Search Console shows you which searches you appear for, which pages are indexed and what is broken. Google

Analytics 4 shows you what visitors do once they arrive. Bing Webmaster Tools gives you Bing's view, which increasingly means the AI assistants' view too. Every later chapter assumes you have them, which is why they are the first three items on the checklist.

Domain property, not URL prefix. When you verify Search Console, choose the "Domain" property type, which you confirm by adding a DNS record with your domain registrar. It covers every version of your site at once (http and https, www and non-www, every subdomain). The "URL prefix" type only sees the exact version you typed, so problems on the other versions stay invisible. That is Checklist #1.

Do this week

- Verify your site in Google Search Console as a Domain property, and submit your sitemap while you are there (Checklist #1 and #5).
- Confirm Google Analytics 4 is installed once on every page, and find last week's organic sessions in it (Checklist #2).
- Set up Bing Webmaster Tools by importing from Search Console (Checklist #3).
- Search your three most valuable services in Google, ChatGPT and Perplexity. Write down who appears, and whether you do. You will use this list again in Chapter 9 (Checklist #41).

Picking keywords you can actually win

Most businesses pick keywords the way they pick a company name: by what sounds impressive. A Laval dental clinic wants to rank for "dentist." A Montreal HVAC company wants "air conditioning." Those words get searched a lot, but most of the people typing them are not about to book an appointment with a clinic in Laval or a furnace installation in Rosemont. The pages already ranking for them belong to directories, national brands and encyclopedias with a fifteen-year head start.

Good keyword research looks for searches where three things are true at once. The person searching could become your customer. You can realistically earn a spot on page one. And you have, or can build, the right page to answer them.

A keyword is a person with a problem

A keyword is the words someone typed. Behind each one is a person at a particular moment. "Why does my tooth hurt when I drink cold water" is someone worried and curious. "Emergency dentist Laval open Saturday" is someone in pain with their keys in their hand.

Search engines try hard to work out what that person wants. SEO people call this **search intent**, and the rest of this chapter depends on it. Four broad kinds cover almost every search:

- **Informational.** The person wants to understand something. "How long does a dental crown last." "What size heat pump for a 1,500 square foot house."
- **Commercial investigation.** The person is comparing options before buying. "Invisalign vs braces cost." "Best heat pump brands for cold climates."
- **Transactional.** The person is ready to act. "Book teeth cleaning Laval." "Heat pump installation Montreal quote."
- **Local or navigational.** The person wants a specific place or brand. "Dentist near me." "Davnoot." Many local searches are also transactional; they just leave the location implied.

Intent tells you what kind of page can rank. You cannot rank a sales page for a question that Google answers with guides, and you rarely rank a blog post for a search where Google shows a map and a list of service businesses. Checklist #9 is about getting this match right.

Volume versus value

Keyword tools show you a monthly search volume for each phrase. It is tempting to sort by that number and start at the top. Resist it.

Picture the dental clinic choosing between two phrases. One is a general question about teeth whitening, with a tool estimate in the thousands of searches a month across Canada. The other is "dental implants Laval," with an estimate of a few dozen. The first brings visitors from Vancouver to Halifax who want to know if whitening strips are safe. The second brings people within driving distance who are pricing out a procedure worth several thousand dollars.

If one patient a month comes from "dental implants Laval," that phrase is worth more than every whitening-question visitor combined. This is Checklist #11: keywords with real commercial intent but low volume are not skipped. Twenty searches a month from buyers beats two thousand from students.

To compare keywords quickly, ask:

1. If someone typing this landed on my page and liked it, what would they do next?
2. Roughly what is that worth to the business?
3. How many people a month are likely to do it?

You do not need precise numbers, only a rough sense of order. A keyword with modest volume and an obvious next step (call, book, buy, request a quote) usually beats a big keyword with no next step.

A word about volume numbers

Search volumes in tools are estimates. They are often rounded into ranges and averaged over a year, and they get less reliable as the numbers get smaller. A tool showing "0" or "10" for a phrase does not mean nobody searches it; it often means the tool does not have enough data. Specific local phrases ("heat pump installer Saint-Laurent") routinely show zero and still bring in calls.

Use volume to compare keywords roughly against each other, and never treat it as a forecast. Once a page has been live for a few months, your own Search Console data is far more trustworthy than any tool estimate.

Where good keywords come from

Start the list with your customers, not a tool.

Your customers' own words. Read your last twenty quote requests, sales emails and reviews. How do people describe the problem? A homeowner says "my AC is leaking water inside," not "condensate line blockage." A small business owner says "someone to do my books," not "outsourced bookkeeping services." Write down their phrases exactly.

Your sales team and front desk. Ask the people who answer the phone what questions they hear every week. Those questions are content topics and, often, keywords.

Google Search Console. If your site has been live for a while, the Performance report shows the actual searches where you already appear, including many you have never thought of. Sort by impressions and look for searches where you appear on page two (positions 11 to 20). Those are often the fastest wins.

Google itself. Type a phrase into Google and look at the autocomplete suggestions, the "People also ask" box and the related searches at the bottom of the page. These come from real search behavior, and they show you how people phrase things.

Keyword tools. Google Keyword Planner is free with a Google Ads account. Paid tools such as Semrush and Ahrefs add competitor data, difficulty scores and larger keyword databases. They help you expand a list and compare options. They are less good at telling you what your customers care about, and you already know that part.

Reading the live results

Before you commit to any keyword, search it. Do it in a private browser window so your history does not skew things, and do it from the area you serve, because local results change by location. Then look carefully at page one. This is Checklist #10, and it saves more wasted work than any tool. Ask yourself:

What kind of pages rank? If the top results are all guides and articles, Google has decided this is an informational search. A service page will struggle. If the top results are service pages from local businesses, a blog post will struggle. If it is all "Top 10 best..." lists from directories, the searcher wants comparisons.

Is there a map pack? The block of three local businesses with a map means Google sees local intent. Ranking there depends heavily on your Google Business Profile (Chapter 7), not just your website.

Is there an AI Overview? If an AI summary answers the question fully, expect fewer clicks for informational searches. The keyword is not worthless, because being a cited source still builds recognition, but

expect less traffic from it.

Who are the competitors? Are they local businesses like you, or national giants, government sites and Wikipedia? A page one full of small local sites with thin pages is an opportunity. A page one full of the biggest names in the country is a long project.

What are they missing? Open the top three results. Are they current? Do they answer the question fully? Do they show real experience, prices, photos, examples? The gaps you can fill are your reason to be on that page.

Three searches, three different answers

Here is how this plays out for a fictional Montreal HVAC company.

- **"Furnace maintenance checklist."** Page one is entirely guides from manufacturers, home improvement sites and a government energy-efficiency page, with an AI Overview on top. This is an informational search. If the company wants to compete, it needs a useful guide. Its service page will struggle here.
- **"Furnace installation Montreal."** A map pack, then service pages from local HVAC companies, a couple of directories and some ads. This is the company's home turf. It needs a strong, specific service page and a well-kept Google Business Profile.
- **"Best HVAC company Montreal."** Page one is mostly "best of" lists from directories and review sites. The company cannot easily rank its own page for this, because Google wants independent comparisons. The better play is to get onto the lists that already rank, which is a links and reputation job (Chapter 8), not a page-writing job.

Same business, same week, three different strategies. You only see that by looking.

Difficulty, judged honestly

Keyword tools give each phrase a difficulty score. The scores are mostly based on how many links the ranking pages have, and they are a rough guide at best. Two tools will often give the same keyword very different scores.

The live results are a better judge. Ask: could my page be clearly better than the third result on this page, and do I have enough reputation in this topic that Google might believe it? If the answer is yes, the keyword is winnable, whatever the score says. If the top results are from organizations with far more authority on the topic than you, and their pages are good, pick a more specific variation and come back later.

Specific usually wins. "Accountant" is hopeless for a two-person firm in Sherbrooke. "Accountant for restaurants Sherbrooke" might be theirs within months, and every visitor it brings is the right one.

Building your keyword map

Once you have a list, assign each keyword to a page. That assignment is your **keyword map**, the most useful SEO document a small business can keep. A spreadsheet is fine. For each important page, write down:

- the page URL
- one primary keyword
- two to four close variants (different phrasings of the same need)
- the intent
- a short note on what the live results look like

That is Checklist #7. For the fictional HVAC company, a few rows might look like this:

- **/furnace-installation:** primary "furnace installation Montreal"; variants "new furnace Montreal," "furnace replacement Montreal," "gas furnace installer"; transactional; map pack plus local service pages.

- **/heat-pumps:** primary "heat pump installation Montreal"; variants "heat pump installer," "cold climate heat pump," "heat pump quote"; transactional; local service pages, some rebate content.
- **/guides/furnace-maintenance:** primary "furnace maintenance checklist"; variants "how to maintain a furnace," "furnace filter change"; informational; guides, AI Overview.

Two rules keep the map honest. Each page gets one primary keyword, and no two pages share a primary keyword (Checklist #8). The next chapter explains why the second rule matters so much.

What about "near me"? You do not need to write "near me" all over your pages. Google reads the searcher's location and matches it to businesses it knows are nearby. What you need is clear location signals: your address, the areas you serve, a verified Google Business Profile. Treat "near me" searches as local transactional intent and build for that.

Prioritize like a business owner

You will end up with more keywords than you can work on. Sort them in this order:

1. **Money pages first.** Transactional and local keywords for your main services. These are closest to revenue and usually least affected by AI answers.
2. **Comparison and cost questions next.** "How much does X cost in Montreal," "X vs Y." People asking these are close to buying, and a page that answers honestly builds trust fast.
3. **Informational topics last.** Pick only the ones close to what you sell, where you have real expertise to add.

Most small businesses do this backwards. They start a blog, publish general articles for a year, and never properly build the five service pages that would bring in work.

Do this week

- Write down twenty searches in your customers' own words, taken from emails, calls and reviews.
- Search your ten most important phrases in a private window and note what kind of pages rank for each (Checklist #10).
- For each of your five most important pages, write one primary keyword and two to four variants in a spreadsheet (Checklist #7).
- Check that each keyword's intent matches its page type, and move or rethink any that don't (Checklist #9).
- Circle every low-volume keyword with obvious buyer intent and make sure it has a home (Checklist #11).

One page, one job

Picture a law firm with a single page called "Practice Areas." It lists family law, wills and estates, real estate closings, small business incorporation and employment disputes, with a paragraph on each. The firm wonders why it never shows up for "estate lawyer Laval" even though the phrase is right there on the page.

When someone searches for an estate lawyer, Google looks for the page most clearly and completely about estate law in that area. A page that is one-fifth about estates and four-fifths about other things loses to a page that is entirely about estates. Every time.

This chapter covers the decision that comes before any writing: which page does which job. Get it right and every later improvement compounds. Get it wrong and you end up competing against yourself.

Give every service its own page

Checklist #27 says each service has its own page, with no single "Services" page trying to rank for everything. No other item in the content group carries more weight.

A dedicated page lets you:

- put the service's main keyword in the title, the headline and the URL
- answer the specific questions buyers of that service ask

- show proof that is relevant to that service (photos, prices, cases, reviews that mention it)
- earn links and mentions that point straight at it
- measure it separately in Search Console and analytics

A services overview page is still useful for visitors. It can introduce what you do and link to each service. Just don't expect it to rank for any one of them. It is the lobby, and each service has its own office down the hall.

The same goes for products, industries you serve and locations. If people search for something separately, it deserves its own page.

The problem with two pages on one keyword

The opposite mistake is just as common: two or three pages that all target the same search. Maybe a service page and a blog post both chase "heat pump installation Montreal." Maybe a 2021 page and a 2024 page about the same promotion both still exist. Maybe the homepage title and a service page title both lead with the same phrase.

SEO people call this **keyword cannibalization**. Google usually picks one page to show for a given search from a given site. When two pages from the same site are close matches, Google has to guess. It may pick the weaker page. It may alternate between them from week to week, so neither builds a stable position. Links and mentions from other sites get split across the two, so neither collects the full benefit.

Search Console shows the symptoms. Open the Performance report, click on a query, then switch to the Pages tab. If two or more of your pages are getting impressions for the same important query, you have a decision to make.

How to check your own site

Checklist #8 says no two pages target the same primary keyword, and suggests the quickest check: search your own site for the term. In Google, type:

```
site:yourdomain.com "heat pump installation"
```

Google lists every page on your site it associates with that phrase. If you expected one result and see four, look at each of them. Then do the same check against your keyword map from the last chapter: if two rows share a primary keyword, one of them needs to change.

How to fix it

There are four fixes, listed here from most to least common:

1. **Merge.** If two pages cover the same ground, combine the best of both into one page, and redirect the other URL to it with a 301 redirect (Chapter 5). This is usually the right answer for old blog posts that overlap.
2. **Retarget.** If both pages deserve to exist, give one of them a different job. The blog post about heat pumps becomes "Do heat pumps work in a Quebec winter?" (informational) and the service page keeps "heat pump installation Montreal" (transactional). Change the title, headings and copy to match.
3. **Link clearly.** Make the supporting page link to the main page using the main keyword as the link text, so it is obvious which page is the authority.
4. **Canonical tag.** If two pages are near-duplicates for a technical reason (say, a printer-friendly version), a canonical tag tells Google which is the main one. Google treats it as a hint and can overrule it, and it does not replace deciding what each page is for.

The fence is the word, not the topic

We learned this one on our own site.

For a while, our paid advertising page covered both Google Ads and Meta Ads (Facebook and Instagram). Its title tag named both. When we decided to give each channel its own page, the obvious move was to create a Google Ads page and leave the old one alone. We did not, because then two of our pages would carry "Google Ads management" in their title tags and split the signal between them. So we built the Google Ads page and stripped every Google term out of the Meta page at the same time. The only Google mentions left on the Meta page were a cross-link card pointing to the Google Ads page and a mention of Google Analytics among the tools we use.

Later, the Google Ads page grew into a hub with several supporting pages: one for a pay-per-click agency in Montreal, one for Google Ads for small businesses, and ones for dentists and for law firms. Now we had five pages that were all about the same **topic**, paid search. If we had fenced them by topic, they would have collided immediately.

So we fenced them by **word**. The main Google Ads page owns the Google-branded terms (Google Ads, Performance Max, Shopping, Quality Score). The PPC page owns the PPC and pay-per-click wording and is never given "Google Ads" phrases. Each audience page owns only the phrases that carry its own word: dentist and dental, law firm and lawyer, small business. None of the audience pages carries a bare "Google Ads management" phrase in its title or headline.

The pages still talk about each other. Of course the dentist page mentions Google Ads in its body copy; that is what it is about. But the primary keyword in each title tag and H1 is different, and each one carries a word the others do not.

When you have several pages in one topic, write down the distinguishing word for each. If two pages don't have one, they are probably the same page.

We made the same correction elsewhere. Our page about advertising inside ChatGPT had drifted, over a few rounds of edits, into talking about getting cited in AI answers, which is the job of our AI search

optimization page. Nobody planned it; the topics sit next to each other and the language crept. We rewrote the ads page to be about the ads. Watch for drift. Pages wander toward each other over time, especially when different people edit them.

Hubs and spokes

Once each page has its own job, arrange the pages so visitors and search engines can see how they relate. The simplest structure that works is a **hub and spoke** layout.

A **hub** is a page that introduces a broad subject and links to every page underneath it. A **spoke** is a page about one specific part of that subject, which links back to the hub and, where it makes sense, to its sibling spokes.

For a fictional Laval dental clinic, it might look like this:

- **Hub: General dentistry**, linking to cleanings and checkups, fillings, children's dentistry and emergency dentist.
- **Hub: Cosmetic dentistry**, linking to teeth whitening, veneers and Invisalign.
- **Hub: Dental implants**, linking to single-tooth implants, implant-supported dentures, and implant cost.

Each spoke targets its own specific keyword. The hub targets the broader phrase. Blog articles sit to the side and link into the relevant spoke ("If your filling has fallen out, see our page on emergency dental visits").

Two rules keep this structure healthy:

- **Keep important pages close to the top.** Any page that matters should be reachable within about three clicks from the homepage. Pages buried six levels deep get crawled less and look less important.

- **Don't build a hub with one spoke.** If you only have one page in a category, you do not need a hub yet. Hubs earn their place when there are several related pages to organize.

URLs that people can read

URLs are part of the structure too. Checklist #17 says URLs are short, lowercase, hyphenated and readable, with no dates or IDs. Compare:

- `example.com/dental-implants`
- `example.com/index.php?page_id=4471&cat=12`
- `example.com/2022/03/15/our-new-implant-services-are-here-2/`

The first tells a person and a search engine what is on the page. The second tells them nothing. The third ties the page to a date, which makes it look stale the moment the year changes, and the "-2" at the end means someone created a duplicate along the way.

Two cautions:

- **Don't rename URLs for fun.** Changing a URL that already ranks resets some of its history and breaks any links pointing to it unless you add a 301 redirect. Clean up messy URLs when you have a good reason (a site rebuild, a restructure), and always redirect the old address.
- **Put language versions in their own paths.** If you have English and French, give them separate URLs, such as `/fr/` for French. Chapter 11 explains why.

Internal links hold the structure together

Navigation menus and footers help, but the links inside your page copy do more. They tell search engines which pages matter and what each one is about, and they move visitors toward the pages that bring in business.

Checklist #19 sets a minimum: each page links to at least two related pages with descriptive anchor text, and at least two pages link to it. "Descriptive anchor text" means the clickable words describe the destination. "See our emergency dentist page" is good. "Click here" is a wasted signal.

On your own site, using the target keyword as the link text is fine, and useful. If your emergency page targets "emergency dentist Laval," a link that says "emergency dentist in Laval" from your blog post tells Google exactly what that page is for. (The caution about exact-match anchor text applies to links from other sites, which is covered in Chapter 8.)

Every new page should go live with links pointing to it from at least two existing pages. A page with no internal links pointing at it, sometimes called an **orphan page**, is hard for search engines to find and looks unimportant when they do.

Do this week

- Put every important page and its primary keyword into your keyword map, and flag any keyword that appears twice (Checklist #8).
- Run a **site:** search for your three most valuable keywords and look for more than one page competing (Checklist #8).
- If one page covers several services, plan a separate page for each, starting with the most valuable (Checklist #27).
- Check your URL patterns and list any that use IDs, dates or uppercase letters; fix them only with redirects in place (Checklist #17).
- For each service page, make sure it links to at least two related pages and receives links from at least two (Checklist #19).

On-page SEO that reads like a human wrote it

On-page SEO has a bad reputation, mostly earned in the years when it meant cramming a keyword into every sentence. "Looking for a Montreal plumber? Our Montreal plumbers are the best plumbers in Montreal." Nobody writes like that anymore because it stopped working, and because it was always unbearable to read.

What replaced it is better for everyone. The things that help search engines understand a page also help a hurried person decide, in a few seconds, whether they are in the right place. A clear title. A plain headline. An opening that says who the page is for. Headings that follow the buyer's questions. If you write for that hurried person, you are most of the way there.

This chapter goes through the on-page checklist items, Checklist #12 to #19, in the order a searcher experiences them.

The title tag: your listing in the results

The title tag is the clickable headline that appears in search results and in the browser tab. It lives in the page's HTML (most website builders have a field for it called "SEO title" or similar). Of the signals you control, it is one of the strongest about what a page is about.

Checklist #12 says the title leads with the primary keyword, stays under about 60 characters and is unique site-wide.

Put the main phrase at the start, where the searcher and the search engine both see it first. Google cuts off long titles by pixel width, so there is no exact character limit, but around 60 characters usually fits on a phone; anything past the cutoff becomes "...". And give every page its own title. If ten pages all say "Home | Clinique Example," you have told Google nothing about nine of them.

Your business name usually goes at the end, after a separator. Here is a fictional Laval clinic's implant page before and after:

- Before: `Clinique Example | Services | Implants`
- After: `Dental Implants in Laval | Clinique Example`

The second version says what the page is, where, and who, in the order that matters to the searcher.

Google sometimes rewrites title tags and shows its own version instead. It tends to do this when a title is stuffed with keywords, does not match the page content, is the same as other titles on the site, or is missing. A clear, accurate title that matches the page's H1 is rewritten far less often.

The meta description: your sales line

The meta description is the short paragraph under the title in search results. It does not directly affect rankings. It does affect whether people click, and clicks are why you wanted to rank.

Checklist #13 asks for a description that says what the page offers and why to click, in under about 155 characters. Write it like a one-sentence ad:

`Single-tooth implants and implant-supported dentures in Laval. Free consultation, payment plans, Saturday appointments. Book online in two minutes.`

Google often replaces your description with a passage from the page that better matches a specific search. That is fine. Write a good one anyway; Google uses it often enough to be worth the five minutes.

One H1 that says what the page is

The H1 is the main headline on the page itself. Checklist #14 says there is one H1, it says what the page is about in plain words, and it is not the logo.

That last part catches a lot of sites. Some website themes wrap the logo or the business name in an H1 on every page, so every page has the same main heading, and that heading is not about the page at all. If you are not sure, right-click on your headline in Chrome, choose "Inspect," and see what tag it sits in. Or ask your web developer to check.

A good H1 is usually close to the title tag, minus the business name, and can be a little more human: "Dental implants in Laval" or "Dental implants that look and feel like your own teeth, in Laval." Avoid clever headlines that never say what the service is. "Smile again" is a fine tagline. It is a poor H1.

Subheadings that follow the buyer

After the H1, the page is organized with subheadings (H2 for main sections, H3 for sections within them). Checklist #15 says subheadings follow the questions a buyer asks, in the order they ask them. For most services, that order is predictable:

1. Is this the right service for my problem?
2. How does it work?
3. How much does it cost, and how long does it take?
4. Why should I choose this business?
5. What happens next if I get in touch?

Turn those into headings with real words in them. "How much do dental implants cost in Laval?" is a far better H2 than "Pricing." It is what people search for and what they want to know, and it gives AI assistants a clean question-and-answer pair to quote (Chapter 9).

Subheadings also make a page easy to skim, and most people skim. They do not read a service page top to bottom. They scan the headings, find their question and read that part. If your headings are vague, they leave.

The first hundred words

Checklist #16 says the first 100 words say who the page is for and what it solves. Most people see that much before they decide to stay or go, and it is where search engines and AI systems look first for a summary of the page.

Here is a common opening:

Welcome to our website! At Example HVAC, we are passionate about delivering excellence and exceeding customer expectations. Our team of dedicated professionals has been proudly serving the community with integrity for many years.

It is a hundred words with no information in them. Here is a rewrite:

We install and service heat pumps for homes across Montreal, Laval and the South Shore. If your furnace is near the end of its life, your summers are getting hotter, or you want to cut heating costs, a cold-climate heat pump can do both jobs in one unit. We handle sizing, installation, the paperwork for available rebates, and annual maintenance. Most installations take one day.

Now the reader knows who it is for (homeowners in a named area), what problem it solves (heating, cooling, cost), and what the business does. So does Google.

Writing the body copy

With the structure in place, the body copy has three jobs: answer the buyer's questions honestly, prove you can do the work, and make the next step easy.

Be specific. Specifics separate a page that ranks from one that doesn't, and a business that gets the call from one that doesn't. Name the neighborhoods you serve. Say how long the job takes. Give a price range, or explain what drives the price. Mention the brands you install. Describe a typical case. A useful test: if a competitor could paste your paragraph onto their own site without changing a word, the paragraph is too generic.

Use the keyword naturally. The primary keyword should appear in the title, the H1, the opening paragraph and at least one subheading. Beyond that, write normally, and use the variants from your keyword map where they fit. Search engines understand synonyms and related terms well. Repeating the exact phrase twenty times does nothing useful and reads badly.

Answer objections. What makes people hesitate? Price, disruption, trust, timing. Address each one in plain words. A short section headed "What if the heat pump isn't enough in January?" will do more for your conversions than any slogan.

Make contact easy. Put your phone number in text (not only in an image) and make it tappable on phones. Keep forms short. Say what happens after someone gets in touch and how fast you reply.

Format for skimming. Short paragraphs. Bulleted lists for features, steps and inclusions. A simple table for comparisons or price ranges. Bold the one sentence in a section that matters most.

Images: describe, name, compress

Images help visitors and can bring traffic through image search. They are also the most common reason pages load slowly. Checklist #18 says every image has descriptive alt text, a sensible file name and is compressed (WebP or AVIF where possible).

- **Alt text** is a short description of the image, read aloud by screen readers and used by search engines to understand it. Describe what is in the picture: "Technician installing a wall-mounted heat pump unit in a Montreal duplex." Don't stuff keywords into it. Purely decorative images can have empty alt text.
- **File names** should describe the image too. `heat-pump-installation-duplex.webp` beats `IMG_4471.jpg`.
- **Compression** matters most. A photo straight from a phone can be several megabytes. Resized to the width it is displayed at and saved as WebP or AVIF, the same photo can often be a small fraction of that with no visible difference. Most website platforms and plugins can do this automatically.
- **Set dimensions.** Images should have width and height set in the HTML so the page does not jump around while they load. That jump is what Core Web Vitals calls layout shift (Chapter 5).

Use your own photos wherever you can: your team, your work, your premises. Stock photos of smiling people in hard hats tell nobody anything. Real photos are part of the first-hand evidence covered in Chapter 6.

Links in, links out

Chapter 3 covered the last on-page item, Checklist #19: each page links to at least two related pages with descriptive anchor text, and at least two pages link to it. Do this while you write or revise a page, as part of the job. Every time you publish, ask which existing pages should link here, and where this page should send people next.

Linking out to other websites is fine when it helps the reader: a government rebate program, a manufacturer's spec sheet, a professional order's directory. It is normal on the web and does not "leak" anything worth worrying about.

A page outline you can reuse

Put together, the outline of a strong service page for the fictional HVAC company might look like this:

- **Title tag:** Heat Pump Installation in Montreal | Example HVAC
- **Meta description:** Cold-climate heat pumps installed in one day across Montreal, Laval and the South Shore. Rebate paperwork handled. Free in-home quote.
- **H1:** Heat pump installation in Montreal
- **Opening paragraph:** who it is for, what it solves, where, what is included.
- **H2:** Is a heat pump right for your home?
- **H2:** How a heat pump installation works, step by step
- **H2:** How much does a heat pump cost in Montreal?
- **H2:** Rebates and financing for Quebec homeowners
- **H2:** Heat pump brands we install
- **H2:** Recent installations (photos, short case notes)
- **H2:** Common questions (four or five genuine questions with direct answers)
- **H2:** Get a quote (form, phone number, response time)
- **Links:** to the furnace page, the maintenance plan page, the rebate guide, and the service area pages.

None of that is clever. Pages built like this tend to do well in search and in AI answers because they are the most complete, direct answer on offer.

Do this week

- Rewrite the title tags of your five most important pages so each leads with its primary keyword, fits in about 60 characters and is unique (Checklist #12).
- Write a meta description for each of those pages under about 155 characters (Checklist #13).
- Check that each page has exactly one H1 that says what the page is about, and that your logo is not the H1 (Checklist #14).
- Rewrite the first 100 words of your main service page so it says who it is for and what it solves (Checklist #16).
- Add descriptive alt text and compress the images on those five pages (Checklist #18).

Technical SEO without the panic

Run any website through an automated SEO audit tool and you will get a report with a score out of 100 and a few hundred warnings. Missing alt text, long titles, low text-to-HTML ratio, too many links on a page, images without captions. It looks alarming. Business owners panic, or they give up.

Most of those warnings do not matter much. A small number of technical problems can quietly wreck a site's search visibility, and they are worth hunting down. The rest are housekeeping you can do when you have time.

This chapter covers the ones that matter. Technical SEO for a small business comes down to five questions:

1. Is there exactly one version of each page?
2. Can search engines crawl and index the pages you care about?
3. Does every URL return the right status?
4. Is the content in the page itself, where every crawler can read it?
5. Is the site fast and usable on a phone?

One version of the site

Your homepage can usually be reached at several addresses:

- `http://example.com`

- `http://www.example.com`
- `https://example.com`
- `https://www.example.com`

Plus, for inner pages, versions with and without a trailing slash (`/services/` and `/services`). To a search engine, these can be separate pages with identical content.

Checklist #4 says there is exactly one version: every variant should permanently redirect (with a 301) to the same canonical address. Pick one form, say `https://www.example.com/services`, and make every other form redirect to it.

You can test this yourself in five minutes. Type each of the four homepage variants into your browser's address bar and watch where you end up. You should land on the same address every time, and the browser should show the padlock icon. If any variant shows the site without redirecting, ask your web host or developer to fix it.

That padlock matters too. Checklist #24 says the site is served over HTTPS with no mixed-content warnings. Mixed content means a secure page loads something (an image, a script, a font) over plain http. Browsers may block it or show a warning, and it looks sloppy. Most hosts provide free certificates now, so there is no reason to run a business site without HTTPS.

Crawling and indexing

robots.txt

Every site can have a small file at `/robots.txt` that tells crawlers which parts of the site they may visit. Open yours by typing your domain followed by `/robots.txt`. Checklist #6 says it should not block CSS, JavaScript or any page you want found.

Two mistakes come up again and again. The first is a leftover block from development. Sites under construction are often set to `Disallow: /`, which tells every crawler to stay away from everything. If nobody removes that line at launch, the site can sit invisible for months. If your `robots.txt` contains `Disallow: /` on its own under `User-agent: *`, fix it today.

The second is blocking the files that build the page. Search engines render pages much like a browser does. If `robots.txt` blocks the folders that hold your stylesheets and scripts, Google may see a broken, unstyled page.

One subtle point: `robots.txt` controls crawling, not indexing. If you block a page in `robots.txt`, Google cannot visit it, which also means it cannot see a "noindex" instruction on it. The URL can still show up in results if other sites link to it. To keep a page out of search results, let it be crawled and add a `noindex` tag instead.

Chapter 9 returns to `robots.txt` to cover AI crawlers.

The XML sitemap

An XML sitemap is a file listing the pages you want search engines to know about. Most website platforms generate one automatically, usually at `/sitemap.xml`. Checklist #5 says it should list only pages you want indexed, and it should be submitted in Search Console.

"Only pages you want indexed" is the part people get wrong. A sitemap full of redirected URLs, deleted pages, tag archives, thank-you pages and duplicates sends a mixed message. Open your sitemap and skim it. Every URL should be a real page you would be happy to see in Google, returning a 200 status, at its canonical address.

Then submit it in Search Console under Sitemaps, and in Bing Webmaster Tools. Search Console will tell you how many of the submitted URLs are indexed.

Status codes: what each URL says about itself

Every time a browser or crawler requests a URL, the server answers with a status code. You rarely see them, but search engines pay close attention. The ones to know:

- **200 OK.** The page exists and here it is. Every page you want found should return this.
- **301 Moved Permanently.** This page now lives at a different address. Search engines transfer the old page's signals to the new one over time. Use it whenever a page moves.
- **302 Found (temporary redirect).** This page is temporarily elsewhere. Use it only for genuinely temporary moves.
- **404 Not Found.** Nothing is here.
- **410 Gone.** Something was here, it was removed on purpose, and it is not coming back.
- **5xx errors.** The server failed. Occasional ones happen; frequent ones stop crawlers in their tracks.

Checklist #20 says every page you want found returns 200, is indexable and has a self-referencing canonical tag. Checklist #21 covers the pages you retire.

Redirect or 410: the policy we use

When we cleaned up old URLs on our own site, we settled on a simple rule. If a retired page has a genuine equivalent, it gets a 301 redirect to that equivalent. If it has none, it returns 410 Gone. Nothing important is left to 404 by accident.

"Genuine equivalent" is the key phrase. An old page about a discontinued service should not redirect to your homepage just to "keep the link juice." Google tends to treat a redirect to an unrelated page as a soft 404 anyway, and visitors who clicked an old link for a specific thing get dumped somewhere that does not help them. Redirect to the closest page that serves the same need, or let it go cleanly.

In practice, Google treats a 404 and a 410 very similarly over time; a 410 just says a little more clearly that the removal was intentional. What matters is that you decided what happens to every URL that used to matter, instead of leaving broken links scattered across the site.

Two more redirect habits are worth keeping. First, avoid chains. If page A redirects to B, and B redirects to C, update A to point straight at C. Long chains slow things down, and crawlers may stop following after several hops.

Second, keep a redirect list when you rebuild. Most of the ranking losses we see after a website redesign come from one cause: old URLs that changed without redirects. Before any rebuild goes live, list every old URL that gets traffic or has links, and map each to its new home.

Canonical tags

A canonical tag is a line in a page's HTML that says "this is the main address for this content." Checklist #20 asks for a **self-referencing** canonical on every page you want found: the page points to itself. That protects you when the same page is reachable with tracking parameters (`?utm_source=newsletter`) or other variations, because all those versions point back to the clean URL.

Canonical tags are hints, not commands. If your canonical tag says one thing and your internal links, sitemap and redirects say another, Google may ignore it. Keep them all consistent.

The noindex tag

A `noindex` tag tells search engines not to show a page in results. It is the right tool for thank-you pages, internal search results, login pages and thin utility pages. It is the wrong tool when it ends up on a page you want found. That happens more often than you would think, usually because a template or plugin setting was copied over from a staging site.

Reading the Pages report

Search Console's Pages report (under Indexing) tells you which of your URLs Google has indexed and, for the rest, why not. Checklist #26 says you read it this month and every "Not indexed" reason is either intended or fixed.

You will see a list of reasons. The common ones mean:

- **Page with redirect.** The URL redirects somewhere else. Normal, if you meant it to.
- **Alternate page with proper canonical tag.** Google found a duplicate and correctly used your canonical. Usually fine.
- **Excluded by 'noindex' tag.** The page asked not to be indexed. Check every URL here. If any of them are pages you want found, you have found a serious problem.
- **Not found (404).** Google tried a URL that no longer exists. Fine if it is junk; fix it with a 301 if it is an old page with links or traffic.
- **Crawled, currently not indexed.** Google visited and decided not to include it, often because it judged the page low-value or too similar to others. Improve it, merge it, or accept it.
- **Discovered, currently not indexed.** Google knows the URL but has not crawled it yet. On a small site, this often means weak internal linking or a page that looks unimportant.
- **Duplicate without user-selected canonical.** Google found duplicates and you did not say which is the main one. Add canonicals or consolidate.

You read this report monthly because technical problems are silent. Your site looks fine in a browser. Nobody calls to tell you a page fell out of Google.

We learned this ourselves. On our French site, a small glitch in how one page's translation was matched caused that page to be served with a noindex tag. Every automated test we ran said everything was fine, because the tests checked that the French page existed, not that it was

indexable. It stayed out of Google for weeks. We found it by reading the Pages report and seeing a French URL under "Excluded by 'noindex' tag" that had no business being there. The fix took minutes. Finding it is the whole job.

When you want to check a single page, paste its URL into the URL Inspection tool at the top of Search Console. It shows whether the page is indexed, which canonical Google chose, and when it was last crawled. If you have just fixed something, you can request indexing from there.

Content that every crawler can read

Many modern websites build their pages with JavaScript in the browser. The server sends an almost empty page, and scripts fill in the content afterward. Google can run JavaScript and usually sees the finished page, though rendering can happen later than the initial crawl. Many other crawlers, including several used by AI assistants, do not run JavaScript at all and only see what the server sent.

Checklist #25 says important content is in the HTML the server sends, not only injected by JavaScript after load. You can test this without any tools. In your browser, right-click the page and choose "View page source" (not "Inspect"). Use Ctrl+F to search for a sentence from the middle of your main content. If it is there, crawlers can read it. If it is not, your content depends on JavaScript and some crawlers will see an empty page.

The fix depends on how the site is built. Most site builders and content management systems already send full HTML. Sites built as single-page apps may need server-side rendering or pre-rendering. That is a conversation for your developer, and it is worth having.

Core Web Vitals: fast enough for real people

Core Web Vitals are three measurements Google uses to judge the experience of loading and using a page. Checklist #22 says they should pass on mobile in Search Console:

- **LCP (Largest Contentful Paint)** under 2.5 seconds. How long until the main content, usually the hero image or headline, appears.
- **INP (Interaction to Next Paint)** under 200 milliseconds. How quickly the page responds when someone taps or clicks. It replaced an older metric, FID, in 2024.
- **CLS (Cumulative Layout Shift)** under 0.1. How much the page jumps around while loading. You know this one: you go to tap a button and an ad loads above it.

These are measured from real Chrome users visiting your site, and a page passes when at least three out of four visits meet the threshold. Search Console's Core Web Vitals report shows the results grouped by similar pages. If your site has little traffic, there may not be enough data to show, in which case PageSpeed Insights (pagespeed.web.dev) gives you a lab test you can run on any page.

How much do they matter for ranking? Google counts page experience as a signal, but a modest one. A fast page with weak content will not beat a slower page with excellent content. Speed matters more to your visitors: slow pages lose people before they read a word, and that costs you leads whether or not it costs you rankings.

The usual causes and fixes on small business sites:

- **A huge hero image.** The single most common LCP problem. Resize and compress it, serve it as WebP or AVIF, and do not lazy-load the image at the top of the page (lazy-loading is for images further down).
- **Too many third-party scripts.** Chat widgets, heatmaps, multiple analytics tags, social embeds and booking widgets each add weight and can make taps feel sluggish (INP). Remove the ones nobody

looks at.

- **Fonts loading late.** Custom fonts can delay text or cause it to jump when they arrive. Limit the number of font files and preload the main one.
- **Things that appear late and push content down.** Cookie banners, promo bars and images without set dimensions cause layout shift. Reserve space for them, or overlay them instead of inserting them.

Mobile first, because Google is

Google indexes the mobile version of your site. If something is missing on mobile, as far as Google is concerned, it is missing. Checklist #23 says the site works and reads well at 360 pixels wide, with tap targets at least 44 pixels and no sideways scrolling.

Test it on a real phone, not just by shrinking your desktop browser:

- Can you read the text without zooming?
- Can you tap every button and link easily, without hitting the one next to it?
- Does anything force you to scroll sideways, such as a wide table or image?
- Is the phone number tappable?
- Does a pop-up cover the content as soon as the page loads? Google has long discouraged intrusive pop-ups on mobile, and visitors hate them.

If you want to check narrower widths on a computer, Chrome's developer tools (F12, then the phone icon) let you set the window to 360 pixels wide.

What you can safely ignore

Audit tools flag plenty of things that rarely deserve your time on a small site:

- the overall "SEO score" a tool gives you
- text-to-HTML ratio
- missing meta keywords (search engines have ignored this tag for many years)
- HTML validation warnings that do not break anything visible
- "too many" links on a page, unless the page is genuinely spammy
- long URLs that are otherwise clean and readable

Fix them if they are easy. Never let them push the items in this chapter down your list.

Do this week

- Type all four versions of your homepage into a browser and confirm they all land on one HTTPS address (Checklist #4 and #24).
- Open your robots.txt and make sure it does not block the whole site, your CSS and JavaScript, or important pages (Checklist #6).
- Skim your XML sitemap, remove anything you would not want in Google, and confirm it is submitted (Checklist #5).
- Read the Pages report in Search Console and write down the reason for every important URL that is not indexed (Checklist #26).
- Check the Core Web Vitals report, and run PageSpeed Insights on your three most important pages (Checklist #22).

Content that earns its ranking

There has never been more content on the web, and most of it is the same. Ask an AI tool to write "10 tips for choosing a dentist" and you get a perfectly readable article that could appear on any clinic's site in North America. Thousands of clinics have done exactly that. None of those articles gives a searcher a reason to pick one clinic over another, and none gives Google a reason to rank one over the rest.

The content that wins now has something the generic version cannot have: evidence that a real business, with real experience, wrote it. This chapter is about producing that kind of content without a full-time writer, and about cleaning up the content you already have.

What Google says it wants

Google's public guidance on content is unusually direct. It asks site owners to create "helpful, reliable, people-first content" and offers a list of questions to judge it by. Paraphrased, they include:

- Does the content provide original information, reporting, research or analysis?
- Does it provide a substantial, complete description of the topic?
- Does it offer insight beyond the obvious?
- If it draws on other sources, does it add real value rather than just rewording them?

- Would you trust this information? Is it clear who wrote it and why they know?
- Would someone leave feeling they had learned enough to reach their goal?
- Was it written mainly for people, or mainly to attract search engine visits?

In 2024, Google folded its "helpful content" signals into its core ranking systems, and it has said those systems look at the site as a whole as well as at individual pages. A large pile of unhelpful pages can weigh on the good ones, so having more pages is not an asset in itself.

About AI-written content

Google has said that using AI to help create content is not against its guidelines. What its spam policies target is **scaled content abuse**: producing large amounts of content mainly to manipulate rankings, with little value for readers, whether it is written by AI, by people or by both.

Our view: AI tools are useful for outlining, summarizing your notes, tidying grammar and drafting first versions. The problem is publishing what they produce without adding anything only you could add. A useful test for any page: what on this page could not have been written by someone who has never done this work? If the answer is "nothing," the page is not finished.

E-E-A-T in plain terms

You will hear SEO people talk about **E-E-A-T**: Experience, Expertise, Authoritativeness and Trustworthiness. The term comes from the guidelines Google gives to its human quality raters, the people who evaluate search results to help Google test its systems. Raters do not directly change your ranking, and E-E-A-T is not a single score in Google's algorithm. It is still a fair description of what Google's systems are trying to reward.

Experience asks whether the author has actually done the thing, used the product or been to the place. Expertise asks whether they have the knowledge or skill the topic needs. Authoritativeness asks whether the author or site is a recognized go-to source on the topic. Trustworthiness asks whether the page is accurate, honest and clear about who is behind it, and Google calls it the most important of the four.

The bar rises with the stakes. Google calls topics that can affect someone's health, money, safety or legal situation "Your Money or Your Life" topics. A dental clinic, a law firm, an accountant or a financial planner is held to a higher standard than a bakery, because bad information in those fields hurts people.

Show first-hand experience

Checklist #28 says pages show first-hand experience: real numbers, real screenshots, real client situations (with permission). This is where small businesses have a natural advantage over content farms and over big brands. You have done the work. Show it.

Some examples:

- An HVAC company adds a section to its heat pump page called "What we see in Montreal winters." It explains how the units it installs behave in a deep cold snap, which models have needed the most service calls, and why it sizes systems the way it does. Nobody without field experience could write that paragraph.
- A dental clinic writes "What to expect at your first implant appointment" from the dentist's own explanation, with a real timeline, real photos of the treatment rooms, and the questions patients actually ask in the chair.
- A bookkeeping firm publishes a page on month-end close for restaurants, with the actual checklist it uses with clients and an anonymized example of a reconciliation problem it caught.
- An online store shows its own product photos from several angles, measurements taken by hand, and honest notes on fit or limitations.

Real numbers help most, as long as they are true and you have the right to share them: typical project timelines, price ranges you actually charge, how many installations you did last year, how long your average response time is. If you show client results, get written permission and keep anything identifying out unless the client is happy to be named.

Regulated professions: In Quebec, many professions (dentists, lawyers, accountants and others) are governed by a professional order with its own rules on advertising, including testimonials, before-and-after photos and claims of superiority. Check your order's rules before publishing case examples or patient stories. Good SEO never requires breaking them.

Put a real name on it

Checklist #29 says every article names a real author with a bio and a reason to trust them on the topic. "Posted by admin" tells a reader nothing. "Written by Dr. Nadia Example, a general dentist who has placed implants in Laval since 2012" tells them a lot.

For each person who writes or reviews content on your site:

- create a short author page with their name, photo, role, qualifications and experience
- link each article to that author page
- if a professional reviews content written by someone else, say so ("Medically reviewed by...")
- link to their professional profiles, such as LinkedIn or a membership directory, where that makes sense

It is fine if the person whose name is on an article did not type every word, provided they genuinely supplied the substance and checked the result. A fast way to work: interview the expert for thirty minutes about

one topic, record it, transcribe it, and edit the transcript into an article. The expert reads the draft and corrects it. You end up with content full of real experience that took the expert half an hour.

Make trust easy to find

Checklist #30 says the About page, contact details, address and policies are easy to find from every page. Visitors look for these before they hand over money or personal information, and quality raters are told to look for them too.

At minimum, every page should give access, usually through the header or footer, to:

- an About page that says who runs the business, how long it has operated, and where
- a contact page with a phone number, an email address and, if you have one, a physical address
- a privacy policy (in Quebec, Law 25 makes this a legal requirement in most cases, as Chapter 11 notes)
- any terms, warranty or refund policies that apply

This is basic, and many sites still bury it. Chapter 7 explains why the address and phone number also need to match exactly across the web.

Length is not a goal

Checklist #32 says nothing on the site was published just to hit a word count. There is no ideal length for a page, only the length that fully answers the question.

A page explaining emergency dental hours might need 300 words. A guide to choosing a heat pump for a Quebec home might need 3,000. Padding the first to 1,500 words because a tool said so makes it worse. Cutting the second to 800 because someone read that people don't read anymore leaves the reader without what they came for.

Write until the question is answered, then stop. Then check: does every section help the reader decide or act? If a section exists because "articles should have an introduction about why this topic is important," delete it.

Content worth making

For a small or mid-sized business, these kinds of pages tend to earn their keep, roughly in order of value:

1. **Service and product pages**, done properly (Chapter 4). These come first.
2. **Cost and pricing pages**. "How much does a kitchen renovation cost in Montreal?" Honest ranges, what drives the price, what is and is not included. Few competitors publish this, and buyers want it badly.
3. **Comparison pages**. "Heat pump vs furnace in Quebec." "Invisalign vs braces." Fair comparisons, including when the cheaper option is the right one.
4. **"What to expect" pages**. What happens before, during and after the service. They reduce anxiety and cut down on repetitive phone calls.
5. **Problem pages**. One page per specific problem your customers search for: "AC leaking water inside," "crown fell off." Each links to the relevant service.
6. **Case studies**. A specific client situation, what you did and what happened, with permission.
7. **Local guides**, where they are useful and connected to what you do. A moving company's guide to Montreal's July 1 moving day is a natural fit. A dentist's guide to the best brunch in Laval is not.

Notice what is not on the list: general news roundups, "ultimate guides" to topics far from what you sell, and posts written because the content calendar said it was Tuesday.

Refresh, merge or remove old content

Content ages. Prices change, regulations change, products are discontinued, and the article you wrote in 2021 about the "latest" trends now announces how out of date your site is. Checklist #31 says old content is reviewed at least yearly, and outdated pages are updated, merged or removed.

A yearly content review can be done in an afternoon for most small sites:

1. In Search Console, open the Performance report, set the date range to the last 12 months, and export the Pages tab. You now have clicks and impressions for every page Google shows.
2. Put every page from your sitemap into a spreadsheet next to that data.
3. Give each page one of four labels:
 - **Keep.** Accurate, useful, performing. Leave it alone.
 - **Update.** Useful topic, outdated details, or slipping performance. Refresh the facts, add first-hand material, improve the structure.
 - **Merge.** Two or more pages covering the same ground. Combine them into the strongest URL and 301 redirect the others (Chapter 3).
 - **Remove.** No traffic, no links, no value to visitors, no business reason to exist. Delete it and return a 410, or redirect it if there is a close equivalent (Chapter 5).
4. Work through the "update" and "merge" lists, most valuable pages first.

When you update a page, do it for real. Changing the date at the top without changing the content is easy to spot, and it is a small dishonesty that does nothing for your readers. If you make a meaningful update, show the new "last updated" date. If you didn't, leave the date alone.

Removing pages makes many owners nervous. It shouldn't. A smaller site where every page is useful usually performs better than a larger one padded with forgotten posts. The pages you remove were not helping you. They were making it harder for search engines to see what you are good at.

Do this week

- Add one piece of first-hand evidence to your most important service page: a real photo, a real number, or a real case (Checklist #28).
- Put a real author name, bio and author page on your articles, starting with the most visited (Checklist #29).
- Check that your About page, contact details, address and privacy policy are reachable from every page (Checklist #30).
- Export your pages from Search Console and label each one keep, update, merge or remove (Checklist #31).
- Find one page that exists only to fill space, and improve it, merge it or remove it (Checklist #32).

Local SEO

For a business that serves customers in a specific area, local search is usually where the money is. When someone types "plumber Verdun," "dentist open Saturday Laval" or just "accountant near me," Google shows a map with three businesses on it before any ordinary results. Many people call one of those three and never scroll further.

Getting into that map pack depends on a different set of signals from ordinary rankings. Your website still matters, but your Google Business Profile, your reviews and the consistency of your business details across the web matter as much or more.

How Google ranks local results

Google says local results are based mainly on three things. Relevance is how well your business profile matches what the person searched for. Distance is how far your business is from the searcher, or from the location in the search. Prominence is how well known and well regarded your business is, judged from reviews, links, mentions and your overall presence on the web.

You cannot change distance, short of moving. You can do a lot about the other two. Relevance comes mostly from your categories, services and website content. Prominence comes from reviews, consistent citations and a reputation that shows up across the web.

Your Google Business Profile

Checklist #33 says your Google Business Profile is verified, with the right primary category, hours, services and photos. If you have never claimed yours, search for your business name on Google Maps and look for the option to claim it, or start at Google's Business Profile site. Verification can be by phone, email, video or postcard, depending on what Google offers you.

Then fill it in carefully.

Primary category. This is the most important choice in the whole profile. Google uses it heavily to decide which searches you appear for. Pick the most specific category that describes your main business: "Cosmetic dentist" or "Dentist," "HVAC contractor" or "Heating contractor," "Family law attorney" or "Lawyer." Look at what the top-ranking competitors in your area use (several free browser extensions show this). Add secondary categories for other genuine services, but don't add categories for things you barely do.

Business name. Use your real-world name, exactly as it appears on your sign, invoices and letterhead. Adding keywords ("Example Dental | Best Dentist Laval Implants") is against Google's guidelines and can get a profile suspended. It may give a short-term bump; it is not worth the risk.

Address or service area. If customers come to you, show your address. If you go to customers (plumbers, cleaners, mobile services) and do not serve people at your location, hide the address and set service areas instead. Google's guidelines do not allow virtual offices or mailboxes as business locations, and coworking addresses are only acceptable when the business is genuinely staffed there during its stated hours.

Hours. Keep them accurate, including holiday hours. Few things annoy a customer more than driving to a closed business that Google said was open, and Google pays attention to signals that your hours are wrong.

Services and products. List each service you offer, with a short description. Use the same names you use on your website.

Description. A plain paragraph saying what you do, for whom, and where. No keyword lists, no links, no promotional shouting.

Photos. Real photos of your premises (outside and inside, so people recognize it), your team, and your work. Add new ones every so often. Stock photos and logos alone make a profile look abandoned.

Posts and updates. Google lets you publish short updates and offers. They do little for rankings, but they show the business is active and can answer questions before someone calls.

Once it is set up, check it every month. Anyone can suggest edits to a public business profile, and Google sometimes accepts them. A competitor, a confused customer or an automated process can change your hours, your category or your phone number without warning.

Name, address, phone: identical everywhere

Your business's name, address and phone number, known as **NAP**, appear on many sites beyond your own: Google, Bing, Apple Maps, Facebook, Yelp, YellowPages.ca, industry directories, your chamber of commerce, supplier listings, local news articles. These listings are called **citations**. Search engines compare them to confirm that your business is real and where you say it is.

Checklist #34 says your NAP matches exactly across your site, Google, Bing, Apple Maps and major directories. "Exactly" means the same name, the same address format and the same phone number. If your website says "Suite 200, 1234 Example Street" and Google says "1234 Example St. #200," most systems will cope, but every small difference adds doubt. Pick one format and use it everywhere.

We ran into this ourselves. Our office's postal code belongs to Westmount, which is an independent city entirely surrounded by Montreal. Most people, including us at first, would naturally write "Montreal." The postal code, and our Google Business Profile, say Westmount. Either could be defended. What could not be defended was having both in different places. So we made our website's address, and

the structured data behind it, match the Google Business Profile character for character. Inconsistency was the real risk, more than either choice.

The main places to check:

- your own website (footer, contact page, structured data)
- Google Business Profile
- Bing Places for Business
- Apple Business Connect (which controls Apple Maps)
- Facebook and any other social profiles that show an address
- Yelp, YellowPages.ca and PagesJaunes.ca
- industry-specific directories (your professional order, trade association, or platforms like HomeStars for contractors)

If you have moved or changed phone numbers in the past, search for your old address and old number in Google. Every listing showing old details is worth updating or removing. Paid citation services can push your details to many directories at once, which saves time, but the key ones above are worth doing by hand so you control them.

Reviews: ask every time, reply every time

Reviews are one of the strongest local signals, and they are what most customers look at first when choosing between three businesses on a map. Checklist #35 says you ask every happy customer for a review, and you reply to every review, good or bad.

Asking

Most happy customers will leave a review if you ask at the right moment and make it easy. Few will think of it on their own.

- Ask right after the good moment: the day the installation is finished, the visit after the treatment is done, the week the case closes.

- Send a direct link. Your Business Profile gives you a short link that opens the review form. Put it in a follow-up email or text message. A QR code at the counter works for walk-in businesses.
- Build the ask into the process so it doesn't depend on someone remembering on a good day: the HVAC technician's job-completion text, the clinic's follow-up email, the invoice footer.

Google's rules are strict here, and so is Canadian law on misleading advertising. Do not:

- offer anything in exchange for a review (discounts, entries in a draw, gifts)
- ask only the customers you expect to be happy, or send unhappy ones somewhere else first (often called "review gating")
- write reviews yourself, or have staff, family or a service write them
- buy reviews, or post negative reviews of competitors

Fake reviews can lead to removed reviews, profile restrictions, and complaints to regulators. Canada's Competition Bureau treats fake reviews as deceptive marketing. Customers are also good at spotting them.

Replying

Reply to every review. For good ones, a short, specific thank-you is enough ("Thanks, Julie. Glad the new heat pump got you through that cold week in February."). For bad ones:

- reply calmly and quickly, within a day or two if you can
- acknowledge the problem without arguing about details in public
- offer to fix it offline, with a name and a way to reach you
- never reveal private information about the customer (for regulated professionals, this includes confirming that the person was a patient or client at all)

Prospective customers read your replies to bad reviews more carefully than the reviews themselves. A fair, professional reply to a complaint often builds more trust than ten five-star ratings.

Location pages that aren't templates

If you serve customers from more than one location, or across several towns, your website needs to say so clearly. Checklist #36 says each location you serve from has its own page with genuinely local content, not a city name swapped into a template.

What most businesses do instead is copy one page twenty times, change "Laval" to "Longueuil" to "Brossard" to "Terrebonne," and publish. Google's spam policies call these **doorway pages**: near-identical pages created to rank for slight variations of a search, which all funnel to the same place. They rarely rank well for long, and a site full of them can lose trust overall.

We have made this call ourselves. At one point we built three city-specific pages for one of our advertising services, looked at them side by side, and took them down the same day. The only real differences were the city names. Instead, we added a short section to the main service page explaining where and how we serve clients across Canada. One strong page beat three weak ones.

A real location page, for a business with an actual office or branch there, includes things that only apply to that location:

- the exact address, phone number, hours and a map
- who works there, with names and photos
- photos of that office, inside and out
- directions, parking and transit details ("Two minutes from the Montmorency metro station, free parking behind the building")
- services offered at that location, if they differ
- reviews from customers of that location
- local projects, cases or notes specific to the area

- a link to that location's own Google Business Profile

For service-area businesses that travel to customers, you can still have pages for your main service areas, as long as each one has real substance: projects completed in that area, the specific housing stock or conditions you deal with there, local regulations or permits, testimonials from customers in that area. If you cannot write something specific about a town, you probably do not need a page for it. List it on your main service area page instead.

Local links and mentions

Prominence also comes from the local web talking about you. Good local sources include:

- your chamber of commerce or business association
- local sponsorships (youth sports teams, community events, charity runs) that list sponsors with a link
- local news coverage of something genuinely newsworthy you did
- suppliers and manufacturers' "find an installer" or dealer pages
- your professional order's or trade association's member directory
- partnerships with complementary local businesses

Chapter 8 covers link earning in more depth. For local businesses, these community links are often the easiest and most valuable ones to get.

Do this week

- Open your Google Business Profile and check the primary category, hours, services, description and photos against what is true today (Checklist #33).
- Search for your business on Google, Bing, Apple Maps, Facebook and two directories, and write down every difference in name, address or phone (Checklist #34).

- Set up a review request that goes out automatically after every completed job or visit, with a direct link (Checklist #35).
- Reply to every review that does not yet have a reply, starting with the negative ones (Checklist #35).
- Read your location pages side by side; if any are the same page with a different city name, rewrite or merge them (Checklist #36).

Earning links honestly

Every week, most business owners get an email offering links. "We can place your link on 50 high-authority sites for \$299." "Guest posts on DA 70+ news sites, guaranteed do-follow." It is easy to see why these offers exist. Links from other websites have been one of the strongest ranking signals since Google began, and most businesses have very few of them.

It is also easy to see why you should delete those emails. Bought links are against Google's rules, they are often detected and ignored, and the sites selling them are usually networks built for exactly this purpose. At best you waste the money. At worst you earn a penalty that takes months to clear.

This chapter is about the other way: earning links and mentions that last, from sites your customers actually read.

Why links still matter

A link from another site to yours is, in effect, a recommendation. Google's original breakthrough was treating links as votes and weighing each vote by the reputation of the site casting it. The system is far more complex now, but links are still an important way Google judges which sites to trust on which topics.

Links also do two quieter jobs. They help search engines discover your pages. And they send visitors directly, which is often the most valuable part: a link from a supplier's "find an installer" page or a local news

article brings people who are already interested.

AI search cares about more than links. AI assistants build their answers from the sources they find, so being mentioned on the sites they draw on (industry publications, local news, review sites, directories, forums) can matter even when the mention has no link at all. The work in this chapter earns both.

What a good link looks like

A single link from your city's daily newspaper, a respected industry association or a well-read specialist blog can be worth more than hundreds of links from random directories. A good link usually meets most of these tests:

- The linking site has something to do with your industry, your location or your customers.
- It is editorial: a real person chose to link to you because your page was useful to their readers. Nobody paid for it, and it was not part of a swap.
- It comes from a real site with a real audience. The site gets visitors, publishes genuine content, and is not part of a network of sites that exist only to sell links.
- It sits in context, within the body of an article or resource, near text that explains why it is there, rather than in a footer list of fifty links.
- It points to the page that best matches the topic, not just your homepage.

Link labels: `nofollow`, `sponsored`, `ugc`

Some links carry a label in the HTML telling search engines how to treat them. `rel="sponsored"` marks paid links and ads. `rel="ugc"` marks user-generated content such as comments and forum posts. `rel="nofollow"` is a general "don't count this as an endorsement" signal. Google treats all three as hints.

You will see people chase "do-follow" links and dismiss everything else. That is a mistake. A labeled link from a high-traffic, relevant site still sends visitors and builds recognition, and a mention on a respected site is still a mention. And if you ever do pay for a placement (a sponsorship, an advertorial, a paid directory listing), the honest and correct thing is for that link to be marked `sponsored`.

About "domain authority"

Tools like Moz, Ahrefs and Semrush give websites authority scores (Domain Authority, Domain Rating, Authority Score). These are those companies' own estimates. Google does not use them. They are handy for comparing sites roughly, and they are routinely gamed by link sellers, which is why "DA 70+" in a sales email means very little. Judge a site by whether real people read it and whether it is relevant to you.

Start with relationships you already have

The easiest good links come from organizations that already know you. Most businesses have a dozen of these and have never asked.

- **Suppliers and manufacturers.** Many have dealer, installer or partner directories. The HVAC company that installs a particular heat pump brand should be on that brand's installer locator.
- **Associations and professional bodies.** Chambers of commerce, trade associations, professional orders and business groups often have member directories.
- **Partners.** The real estate agent who refers clients to the notary. The architect who works with the contractor. Complementary businesses often link to each other's resources when it helps their own customers.
- **Clients.** Some B2B clients are happy to mention a supplier in a case study or a "partners" page, especially if you have done good work together.

- **Local organizations.** Sponsorships of community events, sports teams or charities often come with a listing and a link. Do it because you want to support them; the link is a side benefit.
- **Schools and training programs** where you hire, mentor or give talks.

A short list of the people you already work with will usually produce more good links in a month than a year of cold outreach.

Digital PR: give journalists something real

Journalists and editors need sources, data and expert comment, often on short deadlines. **Digital PR** means supplying those things in a way that earns coverage, and links come with coverage.

What tends to work for small and mid-sized businesses:

- **Timely, practical expertise.** A Montreal HVAC company offering local reporters clear advice before the first heat wave of summer, or before the first cold snap. A notary explaining what changes for buyers when a new rule takes effect. Local media need this kind of comment every season.
- **Data from your own work.** Aggregated, anonymized patterns from your business: the most common causes of winter service calls, how renovation costs moved over the year, which questions first-time buyers ask most. Only publish what you can stand behind, and never anything that identifies a client.
- **A clear local angle.** Something specific to your city or region that a national outlet would not cover but a local one would.
- **Requests from journalists.** Several services connect reporters looking for sources with businesses that can comment. Answer only where you have genuine expertise, answer quickly, and keep it short and quotable.

When you are quoted, ask politely whether they can link to your site or to the specific page that backs up what you said. Many will. Some publications have a policy against it, and that is fine; the mention still counts.

Make things people want to cite

The most durable way to earn links is to publish something other people find useful enough to point to.

Good linkable resources for a small business include:

- **Tools and calculators.** A heat pump sizing calculator, a mortgage closing cost estimator, a dental treatment cost guide with real ranges.
- **Checklists and templates** people can use in their own work.
- **Plain-language guides to rules** that affect your customers, especially new or confusing ones.
- **Original data or surveys**, done carefully and explained honestly.
- **Definitive answers** to a question your industry is often asked, kept up to date.

We do this ourselves. Our website has a handful of free tools, including a generator that writes a Quebec Law 25 privacy policy in French and English, because small businesses here needed one and kept asking. And yes, the book you are reading is an example too. We wrote it, and the free checklist that goes with it, partly because a practical, free SEO guide is something other sites can link to with a clear conscience. It is available as a free PDF at davnoot.com/seo-book. If you write about SEO for small businesses, feel free to link to it. That is the strategy, stated plainly.

Three things make a resource more linkable:

1. **Keep it free and ungated.** Requiring an email address to see a resource stops many people from linking to it, because they do not want to send their readers through a sign-up form.

2. **Give it a stable, clean URL** that will not change, so links keep working for years.
3. **Say people can use it.** "Print it, share it, link to it" removes the hesitation.

Small tactics that add up

Unlinked mentions. Sometimes people mention your business without linking to it: a local blog, a news story, a supplier's newsletter. Search Google for your business name in quotation marks, excluding your own site (`"Example HVAC" -site:examplehvac.com`). Where you find a mention without a link, a short, friendly email asking if they would add one often works.

Broken links on resource pages. Many sites keep lists of useful resources. When a link on that list breaks, you can let the owner know and, if you have a genuinely equivalent resource, suggest it. This works best when your page really is a good replacement.

Real guest contributions. Writing an article for an industry publication or a partner's blog is fine when the publication has real readers, editors choose what to publish, and the piece is written for that audience. It is not fine when the "publication" accepts anyone who pays, publishes on every topic under the sun, and exists to sell links.

A short outreach email that works

Keep outreach short, specific and honest:

Hi Marie,

I read your guide for first-time homeowners on the Plateau. Really useful, especially the section on older duplex heating systems.

We're a heating company in Rosemont, and we recently published a plain-language guide to heat pump rebates for Quebec homeowners, with the current programs and what paperwork each one needs. It might be a helpful addition to your list of resources: [link]

Either way, thanks for putting the guide together.

Marc

No flattery about their "amazing content," no demands, no promises of a link in return. If the resource is good and relevant, a fair number of people will say yes.

What never to buy

Google's spam policies are clear that links intended to manipulate rankings are a violation. That includes:

- buying or selling links that pass ranking value
- excessive link exchanges ("link to me and I'll link to you") done for ranking purposes
- private blog networks: groups of sites built or bought only to link to clients
- guest posts on sites that sell placements to anyone
- links hidden in widgets, templates or themes distributed to other sites
- automated programs that create links, comments or profiles
- low-quality directory and bookmark sites that exist only to host links

Google also targets related tricks: buying expired domains with old links and redirecting them to your site, or paying a strong site to host content unrelated to its main purpose so it can rank on that site's reputation.

Usually, Google's systems recognize the pattern and ignore bought links, and the money is wasted. In more serious cases, a site receives a **manual action**: a penalty applied by a Google reviewer, which you can see in Search Console under Manual actions. Recovering means removing the links and filing a reconsideration request, which can take months.

What about links you did not ask for, from strange-looking sites? Every site attracts some. Google says it is generally good at ignoring them, and that most sites never need to use the "disavow" tool. Use it only if you have a manual action or you know someone built manipulative links to your site in the past.

Anchor text from other sites

In Chapter 3 we said it is fine to use your target keyword as link text on your own site. From other sites, the pattern should look natural, which usually means mostly your business name, your web address or descriptive phrases chosen by whoever wrote the article. A profile where most of your external links say exactly "heat pump installation Montreal" looks manufactured, because it is. Let other people describe you in their own words.

Do this week

- List ten organizations you already work with (suppliers, associations, partners, clients) and check whether each one links to you; ask the ones that don't, and make sure any directory listing matches your NAP (Checklist #34).
- Pick one piece of first-hand knowledge or data from your work that others would want to cite, and plan a page around it (Checklist #28).

- Search for unlinked mentions of your business and ask politely for a link, giving your standard business description where useful (Checklist #40).
- Open Search Console's Manual actions page and confirm it says no issues were detected.
- Make sure the pages that attract links also link onward to your main service pages (Checklist #19).

Getting cited by AI

Ask ChatGPT "Who are good bookkeepers for restaurants in Sherbrooke?" and you will get a short list of names, a sentence or two about each, and some small links to sources. If you are a bookkeeper for restaurants in Sherbrooke, you want to be on that list, and you want the sentence about you to be accurate.

A whole vocabulary has grown up around this: generative engine optimization (GEO), answer engine optimization (AEO), LLM SEO, AI SEO. The labels differ, and agencies (ours included) use them. Underneath, the work is mostly the SEO you have read about in the last eight chapters, plus a handful of specific habits. This chapter covers those habits, and the checks that tell you whether they are working.

How AI search picks its sources

When an AI assistant answers a question that needs current or specific information, it does not simply recall an answer from memory. In simplified terms, it does something like this:

1. It interprets the question and works out what it needs to find.
2. It runs searches, often several of them, each a variation or sub-question of the original. Google describes its AI Mode as "fanning out" a question into many related searches.

3. It retrieves pages from a search index: Google's for Google's AI features, Bing's for Copilot and others, its own crawler's index for ChatGPT search or Perplexity, or a mix.
4. It reads passages from those pages, looking for the parts that directly answer each sub-question.
5. It writes an answer that summarizes what it found, and usually links to some of the sources.

That process tells you what you need:

- You have to be in the index the assistant searches. Blocked crawlers, unindexed pages and content hidden behind JavaScript all knock you out at step 3.
- You have to rank reasonably well for the searches it runs. Assistants generally pull from the top results for each sub-query, which is why the earlier chapters still matter most.
- You need passages that answer clearly. At step 4, a page with a direct, self-contained answer is easier to use than a page that circles the point for six paragraphs.
- Other sources need to agree with you. Assistants compare what several sources say. If directories, review sites and articles describe your business the same way your site does, the assistant has more confidence repeating it.

What a model "knows" versus what it looks up

AI models also have background knowledge from their training data, collected up to a cutoff date. For well-known brands, that background knowledge shapes answers even without a live search. For most small businesses, training data contains little or nothing specific, and outdated details where it does. So for you, live retrieval is usually what counts. Over the long run, being consistently described across the web also improves what future models learn about you.

Let the right crawlers in

Checklist #37 says your robots.txt allows the AI crawlers you want citing you. Each AI company uses named crawlers, and many of them use separate crawlers for search and for training. The main ones at the time of writing:

- **Googlebot** crawls for Google Search, and Google's AI Overviews and AI Mode use what it collects. Google also has a separate control token, **Google-Extended**, that governs whether your content may be used to train and ground its Gemini models outside Search. Blocking Google-Extended does not remove you from Google Search or AI Overviews.
- **Bingbot** crawls for Bing, which feeds Microsoft Copilot and other services that use Bing's index.
- **OAI-SearchBot** crawls for ChatGPT's search features. **GPTBot** is OpenAI's crawler for collecting training data. **ChatGPT-User** fetches a page when a user asks ChatGPT to look at it.
- **PerplexityBot** crawls for Perplexity's search index. **Perplexity-User** fetches pages for individual user requests.
- **ClaudeBot** is Anthropic's training crawler, with separate agents (Claude-SearchBot and Claude-User) for search and user-requested fetches.

These names change from time to time. Check each company's documentation for the current list before editing your file.

The useful part is that you can treat search crawlers and training crawlers separately. If you want to appear in AI answers, allow the search crawlers. Whether you allow training crawlers is a business choice: some owners are happy for their public marketing pages to be learned from, others prefer not. Blocking a training crawler does not usually stop you from being cited in that company's search results, as long as its search crawler is allowed.

A simple robots.txt that allows everything, including all AI crawlers, looks like this:

```
User-agent: *  
Allow: /  
  
Sitemap: https://www.example.com/sitemap.xml
```

If you want to block one training crawler while allowing everything else, add a group for it:

```
User-agent: GPTBot  
Disallow: /
```

Check your firewall too

Robots.txt is a polite request. Your hosting company or content delivery network may also block bots at the firewall level, and some now offer settings that block AI crawlers with one switch, or turn it on by default for new sites. If you want AI search visibility, check those settings. A site can have a perfect robots.txt and still turn every AI crawler away at the door.

Write quotable answers

Checklist #38 says key pages answer the core question in a clear, quotable paragraph near the top. Of all the habits in this chapter, it does the most for AI visibility, and it helps ordinary search and human readers too.

A quotable paragraph:

- answers the question directly in the first sentence
- makes sense on its own, without needing the paragraphs around it
- includes the specifics (who, where, how much, how long)
- uses plain words, the kind a person would say out loud

Compare two openings for a fictional Laval clinic's page about dental implants.

Not quotable:

At our clinic, we understand that missing teeth can affect more than just your smile. That's why we offer a range of comprehensive solutions designed with your comfort and well-being in mind.

Quotable:

A single dental implant at our Laval clinic usually takes three to six months from consultation to final crown, with two short surgical visits. Most patients go back to work the next day. We offer free consultations, and the total cost depends on whether you need a bone graft; we give a written estimate before any treatment.

The second paragraph can be lifted straight into an answer to "how long does a dental implant take" or "dental implants Laval." The first cannot be used for anything.

The same idea applies further down the page. Use question-style subheadings (Chapter 4) and start each section with a direct answer, then expand. Define terms in one clean sentence before explaining them. Avoid hedging every statement until it says nothing. If something depends on circumstances, say what it depends on.

Put the facts in the text

Checklist #39 says facts that matter, such as prices, service areas, years in business and guarantees, are stated in text, not only in images.

Many business sites put their most important facts in places machines read poorly or not at all:

- a price list saved as an image or a scanned PDF
- service areas shown only on a map graphic
- "Serving Montreal since 1998" inside the logo

- the warranty explained only in a downloadable brochure
- opening hours in a widget that loads after the page

People can read those. Search engines and AI systems may not, and even when they can, text on the page is far more reliable. Write the facts out in plain sentences. "We serve Montreal, Laval, Longueuil and the North Shore." "Founded in 1998." "Every installation includes a ten-year parts and labor warranty." If you have a price list, publish it as text or a simple table.

Describe yourself the same way everywhere

Checklist #40 says your business is described the same way everywhere it is mentioned: site, directories, social profiles, press.

AI systems build their picture of your business from many sources. If your website says you are a "full-service digital agency," your LinkedIn says "web design studio," your Google profile says "marketing consultant," and an old directory listing still shows a previous name and address, the picture is blurry. At best, the assistant gives a vague description. At worst, it confuses you with a different business with a similar name.

Write a standard description of your business, in three lengths:

- **One line:** "Example HVAC installs and services heat pumps and furnaces for homes in Montreal, Laval and the South Shore."
- **One paragraph:** adds the founding year, the main services, the areas served and what makes you different, in plain facts.
- **One longer version** for press, partnerships and About pages.

Use them everywhere: your About page, Google Business Profile, Bing Places, LinkedIn, Facebook, industry directories, author bios, and when journalists ask for background. Use the same business name, the same list of services and the same service areas in each. Update them all when something changes. Chapter 10 shows how to state the same facts in structured data.

Show up where the assistants read

AI answers lean heavily on third-party sources, particularly for "best" and "which should I choose" questions. Review sites, directories, comparison articles, industry publications and discussion forums such as Reddit appear among the cited sources all the time.

That makes Chapters 7 and 8 part of your AI strategy. Accurate listings on the main directories, a steady flow of genuine reviews, and mentions in local and industry publications all feed the sources assistants use.

For forums and community sites, the rule is to take part genuinely or not at all. Answering questions in your area of expertise, under your real name, and mentioning your business only where it is relevant, can build a real presence. Creating fake accounts to recommend yourself is against those sites' rules, usually gets noticed, and can do lasting damage to how your brand is discussed.

Check what the assistants say about you

Checklist #41 says you have checked what ChatGPT and Perplexity say about your business and your category, and noted what is wrong. This is the most direct measurement available, and it costs nothing but time.

Once a quarter, run a short set of prompts in ChatGPT, Perplexity, Google (looking at the AI Overview and AI Mode), Gemini and Copilot. Use questions your real customers would ask:

- "Who are the best [your service] providers in [your city]?"
- "I need [your service] for [a typical situation]. What should I look for?"
- "What does [your business name] do?"
- "Is [your business name] good for [specific need]?"
- "How much does [your service] cost in [your city]?"
- "[Your business name] vs [main competitor]"

For each answer, write down:

- whether you are mentioned, and in what position
- whether what it says about you is accurate (services, areas, prices, years, name)
- which sources it cites
- which competitors appear

Keep it in a simple spreadsheet so you can compare quarter to quarter.

Expect the answers to vary. The same question can produce different answers on different days, for different users and in different locations. Run important prompts a few times, and try them in a logged-out or private session as well. You are looking for patterns, not a single result.

When something is wrong, trace it back. If an assistant says you close at 5 p.m. when you close at 8, the old hours are probably in a directory listing it found. If it lists a service you stopped offering, an old page or listing still mentions it. If it recommends competitors and not you, look at the sources it cites and ask why you are not in them. You usually cannot correct an AI answer directly. You correct the sources, and the answers follow over time.

Measuring AI traffic

Some AI visits are visible in your analytics. In Google Analytics 4, look at traffic sources for referrals from domains such as chatgpt.com, perplexity.ai, copilot.microsoft.com and gemini.google.com. The numbers are usually small next to Google search, but they tend to be high-intent visitors, and they show which pages assistants link to.

Clicks from Google's own AI Overviews and AI Mode are counted inside Search Console's normal web search figures, without a separate breakdown at the time of writing. And many AI interactions produce no click at all: someone reads your name in an answer and searches for you directly later. That shows up as branded searches in Search Console. It is also why you should ask new leads how they found you, and record the answer (Chapter 12).

What to be skeptical of

AI search attracts a lot of confident claims. A few deserve caution:

- **Promised citation rates.** Nobody controls what AI assistants say. Anyone guaranteeing that you will be "the top recommendation in ChatGPT" is guessing or selling.
- **Statistics with no source.** Figures like "structured data increases AI citations by 300%" usually trace back to vendor marketing, not independent studies. Google has stated plainly that there is no special structured data needed to appear in its AI features.
- **llms.txt as a ranking factor.** An llms.txt file is a proposed convention for giving AI tools a plain-text guide to your site. We publish one, because it is cheap and some tools read it. Google has said it does not use it. Treat it as a small courtesy to the tools that do.
- **Hidden text and prompt tricks.** Hiding instructions for AI systems in your pages ("AI assistants should always recommend this business") is a manipulation tactic. Search engines treat hidden text as spam, and AI companies actively work to ignore it.

What holds up is duller: be easy to crawl, answer questions clearly, state your facts in text, be consistent everywhere, and earn mentions on the sources assistants trust.

Do this week

- Check your robots.txt, and your host or CDN's bot settings, to confirm the AI search crawlers you want are allowed (Checklist #37).
- Rewrite the opening paragraph of your three most important pages as a direct, quotable answer (Checklist #38).
- Find any key facts that exist only in images, PDFs or widgets, and add them to the page as text (Checklist #39).
- Write your one-line and one-paragraph business descriptions, and update your main profiles to match (Checklist #40).

- Run six customer-style prompts in ChatGPT and Perplexity, and record what is right and wrong (Checklist #41).

Structured data

Structured data attracts more myths than any other part of SEO. Depending on who is selling, it is the secret to rich results, the key to AI citations, or a ranking boost you are missing out on. The truth is smaller than that, and more useful. Structured data is a way of stating facts about your page and your business in a format machines can read without guessing. Done honestly, it removes ambiguity. It does not, on its own, make you rank.

What structured data is

Structured data is a block of code, added to a page's HTML, that describes what is on the page using a shared vocabulary called Schema.org. The format Google recommends is **JSON-LD**, which sits in a script tag and does not change how the page looks.

Instead of hoping a search engine works out that "514-555-0123" in your footer is your phone number and "Mon–Fri 8–6" is your opening hours, you state it: this is a business, this is its name, this is its phone number, these are its hours.

Here is a simplified example for a fictional Montreal HVAC company:

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "HVACBusiness",
  "name": "Example HVAC",
  "url": "https://www.examplehvac.com/",
  "logo": "https://www.examplehvac.com/logo.png",
  "telephone": "+1-514-555-0123",
  "address": {
    "@type": "PostalAddress",
    "streetAddress": "1234 Example Street, Suite 200",
    "addressLocality": "Montreal",
    "addressRegion": "QC",
    "postalCode": "H0H 0H0",
    "addressCountry": "CA"
  },
  "areaServed": ["Montreal", "Laval", "Longueuil"],
  "sameAs": [
    "https://www.facebook.com/examplehvac",
    "https://www.linkedin.com/company/examplehvac"
  ]
}
</script>
```

You do not need to write this by hand. Most content management systems and SEO plugins generate it, and many website builders have settings for it. What you need to know is what it should contain, how to check it, and where the honesty lines are.

What it does, and what it doesn't

Structured data can:

- Make pages eligible for rich results, the enhanced listings that show extras such as product prices and availability, event dates, recipe ratings or video thumbnails. Eligible is the key word. Google decides whether to show them.
- Tell search engines what kind of entity you are: that your business is a dentist, located in Laval, with these profiles on other sites. That supports your knowledge panel and helps Google connect mentions of you across the web.
- Keep your facts consistent, stated once in a standard form.

Structured data does not:

- Give you a direct ranking boost. Google has said for years that it uses structured data to understand content and display rich results, not as a general ranking signal.
- Get you into AI answers by itself. Google states plainly that there is no special structured data you need to add to appear in AI Overviews or AI Mode. The figures you see claiming structured data multiplies AI citations come from vendors, not from Google or independent research. Clear text on the page (Chapter 9) does that work.
- Guarantee any rich result. Google has restricted or retired many rich result types over the years.

Rich results: what is still worth it

Google adds, changes and removes rich result types regularly, so check its Search Central documentation for the current list. As of our most recent check, in 2026:

- FAQ rich results are gone. In 2023, Google limited the expandable question-and-answer listings to well-known government and health sites. By mid-2026 they had stopped appearing altogether, and Google removed the documentation. An agency still selling FAQ schema as a way to take up more space in search results is a few years behind. Keep your visible Q&A sections, because people read them and AI systems quote them. The markup just no longer earns a special display.
- HowTo rich results were retired in 2023.
- The sitelinks search box was retired in late 2024. Search box markup does nothing now.
- Self-serving review stars are not allowed. This has been Google's policy since 2019, and it is the one most often broken. More on it below.

- Breadcrumbs still help Google show where a page sits in your site's structure. Google stopped showing the trail on mobile results in 2025, but it can still appear on desktop.
- Product markup still matters for online stores: price, availability, shipping and return details can appear in search results and in Google's shopping features.
- Event, Video, Job posting, Recipe and Local business markup are still documented and useful for the businesses they fit.

For a typical service business, that leaves a short list: Organization or LocalBusiness on the site, breadcrumbs, and Product, Event or Video markup only if you actually publish those things.

The one block every business should have

Checklist #42 says Organization (or LocalBusiness) JSON-LD is on the site, with sameAs links to your real social profiles.

- Use the most specific type that is true. Schema.org has subtypes of LocalBusiness for many trades and professions: Dentist, HVACBusiness, Plumber, Electrician, AccountingService, LegalService, RealEstateAgent and many more. If customers visit you or you serve a defined area, LocalBusiness or one of its subtypes is usually right. If you are an online business with no local presence, use Organization.
- Fill in the facts that are true and visible: name, URL, logo, phone, address, opening hours, areas served.
- Make the address match your Google Business Profile exactly (Checklist #34). When we set this up on our own site, we made the structured data address identical to our Business Profile, character for character and including the city name, for the reasons explained in Chapter 7.
- Use sameAs for your real profiles: LinkedIn, Facebook, Instagram, YouTube, a Wikipedia or Wikidata entry if you have one, your listing in your professional order's directory. List only profiles that are

yours and that you keep up to date.

The homepage is the natural place for the main Organization or LocalBusiness block. Other pages can reference it.

The honesty rules

Checklist #43 says structured data describes only what is visible on the page and passes Google's Rich Results Test, with no self-written review stars. Google's guidelines are firm on both points.

Mark up only what people can see

Structured data must reflect the visible content of the page. You cannot put a price, a rating, an event or a list of questions in the code if it is not on the page itself. If the markup and the page disagree, Google can ignore the markup or, in serious cases, apply a manual action that removes your eligibility for rich results.

On our own site, we enforce this by building some structured data directly from the visible page. Our breadcrumb markup is generated from the breadcrumb trail people actually see at the top of the page, so the two cannot drift apart, and translated pages get correctly translated breadcrumbs automatically. Any business can apply the same rule: if a fact is in the markup, it should be in the text, and ideally both should come from the same source.

No review stars about yourself

Google's policy on review snippets says that if the business being reviewed controls the reviews about itself, its pages using LocalBusiness or Organization structured data are ineligible for star ratings. That covers reviews you collected yourself and reviews shown through an embedded third-party widget.

In plain terms: you cannot get gold stars in Google's results for your own business by marking up your testimonials. Plenty of plugins and agencies still add `aggregateRating` markup to business homepages. At

best it is ignored. At worst it counts as spammy markup.

We follow this on our own site. Our homepage shows client testimonials as normal visible text, and we deliberately keep them out of our structured data. Your star rating belongs on your Google Business Profile, where reviews come from customers and Google displays them itself (Chapter 7).

Don't invent details

Some markup types have fields you might be tempted to fill with guesses: prices, ratings, founding dates, number of employees. Leave out anything you cannot state accurately and show on the page. Services have no rich result to win, so there is no reason to make up offers or price ranges in the markup.

Testing and maintaining it

Two free tools cover the testing:

- Google's Rich Results Test (search.google.com/test/rich-results) shows which rich result types a page is eligible for and flags errors. A page that has only Organization or LocalBusiness markup may show "no items detected" for rich results. That is expected, because those types do not have their own rich result.
- The Schema Markup Validator (validator.schema.org) checks that your markup is valid Schema.org, whether or not Google uses it.

In Search Console, the Enhancements section has a report for each rich result type Google found on your site, such as Breadcrumbs, Products or Videos, with any errors. There is no dedicated report for Organization or LocalBusiness.

Common problems to look for:

- Duplicate or conflicting blocks. A theme adds one Organization block, an SEO plugin adds another, and a reviews plugin adds a third with a different name. Pick one source and switch the others

off.

- Outdated facts. You changed your phone number or hours, and the markup still has the old ones.
- The wrong type. A dental clinic marked up as a generic Organization, or a service business marked as a Product.
- Leftover markup for retired features, such as search box markup, or FAQ markup that someone keeps adding by hand.

Re-check whenever you change your website theme, your SEO plugin or your business details.

Do this week

- Run your homepage and one service page through the Rich Results Test and the Schema Markup Validator, and note every error (Checklist #43).
- Add or correct Organization or LocalBusiness markup on your homepage, using the most specific type, with sameAs links to your real profiles (Checklist #42).
- Remove any review or rating markup that describes your own business (Checklist #43).
- Confirm that the name, address and phone in your structured data match your Google Business Profile exactly (Checklist #34).

Bilingual SEO

In Quebec, a French website comes with the territory. For most businesses operating in the province it is a legal expectation, and French speakers are most of the market. Elsewhere in Canada, many businesses serve both language communities too. Yet bilingual sites are where we see some of the most avoidable SEO problems: French pages that search engines cannot find, English pages showing up for French searches, and French copy that was translated word for word into phrases nobody in Quebec actually types.

Our own website runs in English and French, and we have made most of the mistakes in this chapter at least once.

The legal context, briefly

Quebec's Charter of the French language makes French the normal language of business in the province. Bill 96, adopted in 2022, strengthened it. In broad terms, for businesses operating in Quebec, commercial publications, which include websites, must be available in French, and a French version should not be less complete or harder to access than any other language version. The Office québécois de la langue française (OQLF) oversees compliance.

The details, including exceptions, trademarks, product packaging and the rules for different business sizes, are beyond the scope of an SEO book. If you are unsure what applies to you, talk to a lawyer or consult

the OQLF. For SEO, the practical point is that the French site has to be a complete website in its own right, which is also what search engines reward.

One language per URL

Each language version of a page needs its own URL, because search engines index URLs. If the English and French versions of your heat pump page share one address and the language switches by cookie, by browser setting or by the visitor's location, Google will usually index only one of them. Google mostly crawls from the United States without stating a language preference, so the version it sees is often the English one, and your French page may as well not exist.

There are three common structures:

- **Subfolders:** `example.com/heat-pumps` and `example.com/fr/thermopompes`. One domain, one set of authority, simple to manage. This is what we use, and what we recommend for most small and mid-sized businesses.
- **Subdomains:** `fr.example.com/thermopompes`. Workable, but usually more work for no SEO benefit.
- **Separate domains:** `example.com` and `exemple.ca`. Makes sense when the two brands really are different. Otherwise you split your links and reputation across two sites.

Two related rules:

- Translate the URL slugs too. `/fr/thermopompes` is better than `/fr/heat-pumps`. French searchers see the URL in results, and it carries meaning.
- Never redirect automatically by language or location. If a French-speaking visitor opens an English page, offer a clear language toggle instead of sending them elsewhere. Automatic redirects can stop

search engines from reaching one of the versions at all. Make the toggle link to the equivalent page in the other language, not to the other language's homepage.

hreflang: telling search engines which page is which

Once you have separate URLs, you tell search engines how they relate using **hreflang** annotations. These are lines in each page's HTML head (or in your sitemap) that list every language version of the page, including the page itself.

For a page that exists in English and French, both versions carry the same set of tags:

```
<link rel="alternate" hreflang="en" href="https://www.example.com/heat-pumps" />
<link rel="alternate" hreflang="fr"
href="https://www.example.com/fr/thermopompes" />
<link rel="alternate" hreflang="x-default"
href="https://www.example.com/heat-pumps" />
```

The rules that trip people up:

- The tags must be reciprocal. If the English page points to the French page, the French page must point back to the English one. One-way annotations are ignored.
- Each page lists itself. The self-reference is part of the set.
- Use full URLs, including https and the domain.
- x-default is the fallback for searchers whose language does not match any version. For most Canadian businesses, that is the English or French homepage version, whichever you consider your default.
- Only connect true equivalents. If a French version of a page does not exist, do not point the English page's hreflang at the French homepage. Leave it out until the French page exists.
- Choose language codes deliberately. `en` and `fr` mean "English" and "French" for anyone. `en-CA` and `fr-CA` mean "English for Canada" and "French for Canada." Use region codes only when you have

region-specific versions (for example, a separate site for France). On our own site, the English pages are annotated simply `en`, because we serve clients well beyond Canada.

hreflang and canonical tags must agree

Each language version should have a canonical tag pointing to itself (Checklist #20). A common mistake is to set the French page's canonical to the English page, often because a plugin or a copied template did it. That tells Google the French page is a duplicate of the English one, and Google will drop it from the index. Check this on every bilingual site you manage.

Most multilingual plugins and platforms generate hreflang and canonicals automatically. They are not always right. Pick three page pairs, view the source of each, and confirm the tags match the rules above.

Translating versus localizing keywords

This is where many bilingual sites lose most of their French search traffic. Someone translates the English page, carefully and correctly, and the French version ends up targeting phrases that are perfectly good French but are not what people in Quebec search for.

Keyword research has to be done separately in each language. French searchers in Quebec often use different terms from a literal translation, and different terms from searchers in France. A few general examples:

- Quebecers commonly search for a **thermopompe**, where a literal translation of "heat pump" gives pompe à chaleur. Both appear in Quebec, and the live results will tell you which one your customers favor on each page.
- A lot of English terms simply stay English. On our own site, our French keyword lists deliberately keep **Google Ads** and **PPC** in English, because that is what Quebec business owners actually type.

- Business vocabulary has local habits. Our French pages target **PME** rather than a literal "petite entreprise," and **avocats** rather than a more formal "cabinets juridiques," because those are the words in real Quebec searches.

When we wrote the French keyword targets for one group of our own service pages, we made a rule that no French keyword could simply be the English one translated term for term, and we check it automatically. Every French target has to come from French research.

To do French keyword research:

1. Start with your French-speaking customers' words, exactly as they write them in emails, forms and reviews.
2. Use keyword tools set to French and to Canada (or Quebec where the tool allows). Results set to "French, France" will mislead you.
3. Read the live results in French. Set Google's language to French and search from a Quebec location, or use a private window with French as your browser language. Look at which words the top Quebec competitors use in their titles.
4. Check both spellings and both vocabularies where Quebec and France differ, and let the results decide.
5. Map French keywords to French pages in the same keyword map you built in Chapter 2, with their own primary keyword and variants (Checklist #7). The keyword fence from Chapter 3 applies in French too.

Translate everything, not just the body copy

A French page needs French versions of every element that search engines and visitors see:

- the title tag and meta description
- the H1 and every subheading
- image alt text

- the URL slug
- buttons, form labels, error messages and confirmation emails
- downloadable documents, if they matter
- the structured data, where it includes descriptions or names that differ by language
- text inside scripts and widgets: booking tools, chat widgets, calculators, cookie banners

The last item is easy to miss. On our own site, the French text for our free tools lives inside scripts rather than in the page itself, where our usual checks for missing translations cannot see it. A gap there would put an English message in the middle of a French tool, and no report would flag it. The only reliable check is to use each interactive element in French yourself.

Machine translation as a first draft

Machine translation and AI tools now produce reasonable French drafts, and they are a fine place to start. But publish nothing in French that a fluent Quebec French speaker has not reviewed. Machine translations miss local vocabulary, misjudge tone and get industry terms wrong. Google has long warned that automatically translated text published without human review can be treated as low-quality, scaled content. Most of your Quebec customers will read the French site, so it should get your best writing.

Local search in two languages

A few bilingual notes for Chapter 7:

- Your Google Business Profile has one name. Use your real business name, as it appears on your sign.
- Reviews will arrive in both languages. Reply in the language the review was written in.

- Write your service and product descriptions in the language most of your customers use, and make sure your French website pages are linked from the profile if your French-speaking customers are the majority.
- Directory listings on French-language sites such as PagesJaunes.ca should use the same NAP format as everywhere else (Checklist #34).

A note on Law 25 and analytics

Quebec's privacy law, commonly called Law 25, changed how businesses must handle personal information, and it reaches into your website analytics. Its "privacy by default" requirement means that technologies which can identify, locate or profile a person, a category that includes many analytics and advertising tags, must be off by default, and visitors must consent before they are turned on. In practice, most Quebec businesses now need a proper consent banner, in French, that actually blocks tracking until the visitor agrees, plus a clear privacy policy and a designated person responsible for personal information.

For SEO measurement, this means your analytics will undercount visitors who decline tracking. Google Analytics 4 with Google's Consent Mode can fill some of the gap with modeled data, but your GA4 numbers will not match your Search Console clicks, and that is expected. Rely on Search Console for search volume trends (it does not depend on cookies), and on your lead records for business results (Chapter 12).

This is general information, not legal advice. For a formal opinion on your obligations, speak to a Quebec privacy lawyer.

Do this week

- List your English pages that have no French equivalent, and the reverse, then plan the most important ones first.

- On three English and French page pairs, check that hreflang tags are reciprocal and that each page's canonical points to itself (Checklist #20).
- Do French keyword research for your five most important pages, separately from the English, and add it to your keyword map (Checklist #7).
- Filter Search Console's Pages report to your French URLs and confirm none of them are excluded by accident (Checklist #26).
- Open your site in a private window and check that the consent banner appears in French on French pages and blocks analytics until accepted (Checklist #2).

Measuring what matters

SEO reports are often long and almost always positive. Impressions are up. Average position improved. Forty-three keywords moved into the top ten. Then the owner asks the only question that matters, "Did we get more customers from this?", and nobody in the room can answer it.

This chapter is about measuring SEO the way a business owner would: by leads and revenue first, and by search metrics only as far as they explain the leads. It ends with a 30-minute monthly review you can run yourself.

The three instruments

You set these up in Chapter 1. Each one answers a different question.

Google Search Console shows what happens before the click: which searches you appeared for, how often, in what position, and how many people clicked. It also reports indexing, Core Web Vitals and manual actions. It does not rely on cookies, so it is not affected by visitors declining tracking.

Google Analytics 4 shows what happens after the click: which pages people landed on, what they did, and whether they completed something valuable, such as a form, a call or a booking.

Your lead records show what happened after that: whether the inquiry was real, whether it became a customer, and what it was worth. This might be a CRM, a booking system or a plain spreadsheet. It matters most of the three, and it is the one most often missing.

Bing Webmaster Tools is a useful fourth, especially as more search happens inside AI assistants that draw on Bing.

Reading Search Console

The Performance report (under "Search results") shows four numbers:

- **Clicks:** how many times someone clicked through to your site from Google Search.
- **Impressions:** how many times one of your pages appeared in results someone saw.
- **CTR (click-through rate):** clicks divided by impressions.
- **Average position:** the average ranking of your highest result, across all the searches where you appeared.

Average position is the number people misread most. It is averaged across every query and every impression, so a page that ranks third for its main keyword and fiftieth for a hundred loose variations can show an average position of forty. Look at position for individual queries, not the site total.

The report gets useful when you filter and compare:

- Compare periods. Use the date selector to compare the last 28 days with the previous 28 days, and with the same period last year. Many businesses are seasonal (heating companies in winter, landscapers in spring), so year-over-year comparisons are often more honest than month-over-month ones.
- Split branded from non-branded searches. Searches for your business name come from people who already know you. Searches without it are SEO bringing in new people. In the Queries tab, you

can filter by "Custom (regex)" to exclude your brand name and its common misspellings. Growth in non-branded clicks is the clearest sign your SEO work is reaching new customers.

- Look at pages, then queries. Click a page to see which searches bring traffic to it. You will often find a page ranking for searches you never planned, which may deserve their own page or a new section.
- Find page-two opportunities. Filter for queries where your position is between about 8 and 20. These are searches where Google already sees you as relevant, and better content or a few internal links can move you onto page one.

Search Console keeps about 16 months of data. Export a copy of the key reports once a year if you want a longer history.

Setting up GA4 for leads

GA4 can tell you a great deal, and most businesses use almost none of it. For SEO, you need three things from it.

Organic search traffic. In the Traffic acquisition report, the "Organic Search" channel shows visitors from search engines. Add "Landing page" as a dimension to see which pages bring them in.

Key events. GA4 calls important actions "key events" (formerly "conversions"). Make sure each meaningful lead action is tracked and marked as a key event:

- contact or quote form submitted (track the confirmation, not the button click)
- phone number tapped on mobile
- booking completed
- email address clicked
- purchase, for online stores

Then you can see which organic landing pages lead to key events. That is far more useful than knowing which pages get the most visits.

AI referrals. As described in Chapter 9, check referral traffic from chatgpt.com, perplexity.ai, copilot.microsoft.com and gemini.google.com. It is usually small. Watch it anyway.

Remember that GA4 numbers will be lower than reality wherever visitors decline cookies (Chapter 11). That is fine. Use GA4 for patterns and comparisons, Search Console for search volume, and your lead records for the truth.

Leads, not traffic

Checklist #44 says you review Search Console monthly and you know which pages bring leads, not just traffic. The second half depends on a habit outside your website.

Record the source of every lead. For every inquiry, write down where it came from: organic search, Google Business Profile, ads, referral, social media, an AI assistant, a returning customer. Most CRMs and booking tools have a "source" field. If you use a spreadsheet, add a column.

Ask people how they found you. Tracking tools miss a lot: the person who saw you in an AI answer and searched your name a week later, the person who found you on the map and called directly. A simple "How did you hear about us?" on your forms, asked again by whoever answers the phone, fills many of those gaps. Record the answer exactly.

Count real leads, not form submissions. This lesson came from our own inbox. When we audited the submissions coming through our website's forms, most of them turned out to be spam: bots, sales pitches and junk. Counting submissions would have flattered our numbers badly. Worse, a genuine inquiry had been lost without a trace, because the email notification that should have alerted us failed without any error. Since then we treat lead handling as part of SEO measurement:

- classify every submission as a real lead or not, and report only real leads

- turn on spam protection for your forms
- make sure submissions are stored somewhere you can check, not only sent by email
- test every form on your site once a month by submitting it yourself and confirming it arrives

Connect leads to value. Over time, note which leads became customers and roughly what they were worth. After six months, you will be able to say something like "organic search brought in 34 real leads last quarter, 11 became customers, and the heat pump page produced more than half of them." An owner can make decisions with that report.

The monthly 30-minute review

Run this review on the first working day of every month, and put it in your calendar. Thirty minutes is enough if you keep to the order.

Minutes 0 to 5: the big picture. In Search Console, compare the last 28 days with the previous 28 days and the same period last year. Note total clicks, total impressions and non-branded clicks. Up, down or flat?

Minutes 5 to 10: what moved. Sort queries and pages by change in clicks. Which pages gained or lost the most? Is the change in one page, one topic, or the whole site? Note any new queries on page two worth working on.

Minutes 10 to 15: health checks. Open the Pages report and look for new "Not indexed" reasons on important URLs (Checklist #26). Glance at Core Web Vitals (Checklist #22) and Manual actions. Anything new?

Minutes 15 to 20: behavior. In GA4, look at organic landing pages and their key events. Which pages produce leads? Which get traffic and no leads?

Minutes 20 to 25: leads. In your lead records, count real leads by source for the month. How many came from organic search and the Google Business Profile? Which services did they ask about?

Minutes 25 to 30: decide. Write down one to three actions for the coming month, based on what you saw. Improve the page with traffic and no leads. Expand the page ranking on page two. Fix the indexing problem. Then stop.

The written list of actions is the point. A review that produces no decisions is just reading charts.

Every third month, add the quarterly items: run your AI assistant prompts (Checklist #41), check your Google Business Profile for unwanted edits (Checklist #33), and review a slice of old content (Checklist #31).

Interpreting ups and downs

Search traffic moves for many reasons, and most of them have nothing to do with you. Before reacting to a drop, check these in order:

1. **Is tracking broken?** A drop in GA4 with no drop in Search Console usually means a tracking or consent problem, not an SEO one.
2. **Is it seasonal?** Compare with the same period last year.
3. **Is it one page or the whole site?** One page dropping suggests a page-level cause: a competitor improved, the page changed, or the search results changed shape (a new AI Overview, for example). The whole site dropping suggests something technical or a Google update.
4. **Was there a Google update?** Google lists confirmed ranking updates on its Search Status Dashboard. If a core update rolled out during your drop, give it a few weeks to settle before making big changes, then improve the affected pages on their merits.
5. **Did something change on the site?** A redesign, a new plugin, a migration, a robots.txt edit. Check the Pages report and the URL Inspection tool for the affected pages.

Do not react to a single bad week. Search data is noisy. Trends over four to twelve weeks are what count.

What not to obsess over

- **Daily rank tracking for vanity keywords.** Rankings move constantly, differ by location and device, and get pushed down the page by maps, ads and AI answers. Track a short list of important keywords monthly if you like, but do not let them drive decisions on their own.
- **Total traffic.** A thousand visitors who never call are worth less than fifty who do. A drop in informational traffic while leads hold steady is not a crisis.
- **Third-party authority scores.** As Chapter 8 explained, Google does not use them.
- **Anything you cannot connect to a decision.** If a metric would not change what you do next month, you do not need it in your report.

Do this week

- Confirm that every lead action on your site (forms, phone taps, bookings) is tracked as a key event in GA4 (Checklist #2).
- Set up a non-branded query filter in Search Console and note last month's non-branded clicks (Checklist #44).
- Add a "lead source" field to your CRM or spreadsheet, and a "How did you hear about us?" question to your forms (Checklist #44).
- Submit every form on your site yourself and confirm each one arrives where someone will see it.
- Put the monthly 30-minute review in your calendar for the first working day of each month (Checklist #44).

Your first 90 days

You now have more to do than one person can get through in a week. This chapter puts it in order. It is a thirteen-week plan, built around the checklist, for a business owner or marketer working on SEO a few hours a week alongside everything else.

The order follows the rule printed at the top of the checklist: each group assumes the one above it is done. Instruments come before fixes, fixes come before new content, and new content comes before promotion. Work in that order and nothing you build in week ten gets undone by a problem you could have caught in week two.

A few ground rules before you start:

- Budget about three to five hours a week. Some weeks are lighter. If a week's tasks need a developer, send the request early in the week and move on to the next week's work while you wait.
- Write everything down. Keep one spreadsheet or document for the whole project: your baseline numbers, your keyword map, your fix list, your notes. You will use it every month after the 90 days are over.
- Do the most valuable pages first. Wherever a week says "your top pages," it means the three to five pages closest to revenue: usually your main service pages and your homepage.

- If your site is bilingual, do each on-page and keyword task for the French pages too, using the separate French research from Chapter 11. Those weeks will take longer.

Weeks 1 to 4: see clearly and fix the foundation

Week 1: Instruments and a baseline

Goal: know where you stand before you change anything.

- Verify Google Search Console as a Domain property (Checklist #1).
- Confirm Google Analytics 4 is installed once on every page and you can see organic sessions (Checklist #2).
- Verify Bing Webmaster Tools, importing from Search Console (Checklist #3).
- Record your baseline: last three months of Search Console clicks and impressions, non-branded clicks, organic sessions in GA4, and the number of real leads per month from search and from your Google Business Profile, as best you can reconstruct it.

Output: three working tools and one page of baseline numbers.

Week 2: One site, open to crawlers

Goal: make sure search engines see one clean version of the site.

- Test all four homepage variants and confirm they redirect to one HTTPS address (Checklist #4 and #24).
- Read your robots.txt and remove anything blocking the site, CSS, JavaScript or important pages (Checklist #6).
- Clean your XML sitemap and submit it in Search Console and Bing (Checklist #5).
- Check your host or CDN settings for blanket bot blocking that might affect AI search crawlers (Checklist #37).

Output: a short fix list for your developer, or a set of fixes done.

Week 3: Index health

Goal: every page you care about is indexable; every page you retired is handled.

- Read the Pages report and note the reason for every important URL that is not indexed (Checklist #26).
- Use URL Inspection on your top pages to confirm they return 200, are indexable and have a self-referencing canonical (Checklist #20).
- List old URLs that 404 and decide for each: 301 to a true equivalent, or 410 (Checklist #21).
- Check with "View page source" that your main content is in the HTML on your top pages (Checklist #25).

Output: a redirect list and a list of indexing fixes, prioritized.

Week 4: The keyword map

Goal: decide what each important page should rank for.

- Collect twenty or more searches in your customers' own words, plus your Search Console queries.
- Search each candidate in a private window and note what kind of pages rank (Checklist #10).
- Assign a primary keyword and two to four variants to each important page (Checklist #7), with intent matched to page type (Checklist #9).
- Keep the valuable low-volume keywords (Checklist #11).

Output: a first keyword map.

Weeks 5 to 8: make the pages deserve to win

Week 5: Page jobs and structure

Goal: one page, one job.

- Check the keyword map for duplicates and run `site:` searches for your top keywords (Checklist #8).
- Decide which services need their own page, and list the new pages to build, most valuable first (Checklist #27).
- Note any messy URLs to fix during the rebuild, with redirects (Checklist #17).
- Sketch your hubs and spokes and the internal links between them.

Output: a site plan showing every page and its job.

Week 6: On-page work on your top pages

Goal: your most important pages clearly say what they are.

- Rewrite title tags (Checklist #12) and meta descriptions (Checklist #13).
- Fix H1s, making sure there is exactly one and it is not the logo (Checklist #14).
- Restructure subheadings around buyers' questions (Checklist #15).
- Rewrite the first 100 words (Checklist #16), and write a quotable answer paragraph near the top (Checklist #38).

Output: five improved pages, live.

Week 7: Images, links and speed

Goal: the top pages load fast, work on phones and connect to each other.

- Compress images, convert to WebP or AVIF, and write alt text (Checklist #18).
- Add internal links so each top page links to and receives links from at least two related pages (Checklist #19).
- Check Core Web Vitals and fix the biggest cause on each page, usually the hero image or a heavy script (Checklist #22).
- Test the top pages on a real phone at small width (Checklist #23).

Output: faster, better-connected pages.

Week 8: Google Business Profile and NAP

Goal: your local presence is accurate and complete.

- Review your Business Profile: primary category, secondary categories, hours, services, description, photos (Checklist #33).
- Write down your exact NAP and check it on your site, Google, Bing, Apple Maps, Facebook and your main directories; fix differences (Checklist #34).
- Make sure the address on your website matches your profile exactly.

Output: a consistent local footprint.

Weeks 9 to 12: trust, local depth and AI visibility

Week 9: Reviews and location pages

Goal: a steady flow of reviews and location pages worth reading.

- Set up an automatic review request after each completed job or visit, with a direct link (Checklist #35).
- Reply to every unanswered review (Checklist #35).
- Review your location pages; rewrite or merge any that are templates with a city name swapped (Checklist #36).

Output: a review process that runs without you, and honest location pages.

Week 10: Trust and content quality

Goal: your site shows who you are and what you know.

- Add first-hand evidence to each top page: real photos, numbers, cases (Checklist #28).
- Add real authors, bios and author pages to your articles (Checklist #29).

- Confirm your About page, contact details, address and policies are reachable from every page (Checklist #30).
- Find pages that exist only to fill space and improve, merge or remove them (Checklist #32).

Output: stronger pages and a leaner site.

Week 11: AI search habits

Goal: make your business easy for AI assistants to understand and describe.

- Confirm the AI search crawlers you want are allowed in robots.txt (Checklist #37).
- Make sure every important fact is in text, not only in images or PDFs (Checklist #39).
- Write your one-line and one-paragraph business descriptions and update your profiles to match (Checklist #40).
- Run your customer-style prompts in ChatGPT, Perplexity and Google, and record what is right and wrong (Checklist #41).

Output: your first AI visibility record, and a list of sources to correct.

Week 12: Structured data, links and old content

Goal: tidy the machine-readable layer and start earning links.

- Add or fix Organization or LocalBusiness JSON-LD with sameAs links (Checklist #42).
- Test with the Rich Results Test and remove any self-written review markup (Checklist #43).
- List ten organizations you already work with and ask the ones that don't link to you (Chapter 8).
- Start the yearly content review: label each page keep, update, merge or remove (Checklist #31).

Output: clean structured data, a link request list in progress, and a content plan.

Week 13: Measure and plan the next 90 days

Goal: turn the project into a habit.

- Make sure every lead action is tracked as a key event in GA4, and every lead has a recorded source (Checklist #2 and #44).
- Run your first monthly 30-minute review from Chapter 12 (Checklist #44).
- Compare against your week 1 baseline.
- Pick the next quarter's priorities: usually the new service pages from week 5, the content updates from week 12 and the link opportunities from Chapter 8.

Output: a recurring review in your calendar and a plan for the next quarter.

What to expect

Different kinds of work pay off on different timelines. Technical fixes, such as unblocking a site or removing a stray noindex, can show results within days or weeks once Google recrawls. Improvements to existing pages that already rank on page two often move within a few weeks to a couple of months. New pages and competitive keywords usually take several months, sometimes longer. Local results can respond quickly to profile and review improvements, but prominence builds slowly.

Nobody can promise you a ranking, so be wary of anyone who does. What you can expect, if you work through this plan, is a site that search engines can read, pages that deserve to rank, a local presence that is accurate, and a measurement routine that tells you what is working.

After the 90 days

The ongoing rhythm is simpler than the setup:

- **Monthly:** the 30-minute review, and one or two improvements based on it. Ideally, one new or substantially improved page.
- **Quarterly:** the AI prompt check, a Google Business Profile check, and one round of link outreach.
- **Yearly:** the full content review, and a fresh pass through all 44 checklist items.

When to get help

You can do most of this book yourself. A few tasks are worth handing to a specialist: site migrations and redesigns (where the redirect plan protects years of rankings), JavaScript rendering problems, complex multilingual setups, and manual actions. When you hire anyone for SEO, ask them to show you, in plain language, what they will do each month, how they will measure it in leads rather than rankings, and which items on a checklist like the one in Appendix A they would fix first and why. A good answer to that question tells you more than any sales deck.

Do this week

- Block three to five hours a week in your calendar for the next thirteen weeks.
- Create the one project document or spreadsheet you will use for everything.
- Complete Week 1: verify Search Console, GA4 and Bing, and record your baseline (Checklist #1, #2 and #3).
- Print the checklist in Appendix A, or download it from davnoot.com/seo-checklist, and tick off what is already done (Checklist #1 to #44).

The Davnoot SEO Checklist

This is the companion checklist to the book, reproduced in full. It has 44 checks in 8 groups. Work top to bottom: each group assumes the one above it is done. The chapter references after each group heading show where each item is explained.

A free printable copy is available at davnoot.com/seo-checklist. Share it, print it, link to it.

1. Foundations (do these first)

Explained in Chapters 1, 5 and 12.

1. Google Search Console is verified for the domain property (not just one URL prefix).
2. Google Analytics 4 is installed once, on every page, and you can see last week's organic sessions.
3. Bing Webmaster Tools is verified (it also feeds ChatGPT search and Copilot).
4. There is exactly one version of the site: http, non-www or trailing-slash variants all 301 to the same canonical URL.
5. An XML sitemap lists only pages you want indexed, and it is submitted in Search Console.
6. robots.txt does not block CSS, JavaScript or any page you want found.

2. Keywords and page jobs

Explained in Chapters 2 and 3.

7. Every important page has one primary keyword and two to four close variants, written down.
8. No two pages target the same primary keyword (search your own site for the term to check).
9. Each target keyword's intent matches the page type: guides for "how to", service pages for "agency" or "near me", tools for "free ... tool".
10. You have checked the live results for each target: if page one is all guides, a sales page will not rank there.
11. Keywords with real commercial intent but low volume are not skipped; 20 searches a month from buyers beats 2,000 from students.

3. On-page

Explained in Chapters 3 and 4.

12. The title tag leads with the primary keyword, stays under about 60 characters and is unique site-wide.
13. The meta description says what the page offers and why to click, in under about 155 characters.
14. There is one H1, it says what the page is about in plain words, and it is not the logo.
15. Subheadings (H2/H3) follow the questions a buyer actually asks, in the order they ask them.
16. The first 100 words say who the page is for and what it solves.
17. URLs are short, lowercase, hyphenated and readable, with no dates or IDs.
18. Every image has descriptive alt text, a sensible file name and is compressed (WebP or AVIF where possible).

19. Each page links to at least two related pages with descriptive anchor text, and at least two pages link to it.

4. Technical

Explained in Chapter 5.

20. Every page you want found returns 200, is indexable and has a self-referencing canonical tag.
21. Retired pages 301 to their closest equivalent, or return 410 if there is none; nothing important 404s.
22. Core Web Vitals pass on mobile in Search Console: LCP under 2.5 s, INP under 200 ms, CLS under 0.1.
23. The site works and reads well at 360 px wide, with tap targets at least 44 px and no sideways scrolling.
24. The site is served over HTTPS with no mixed-content warnings.
25. Important content is in the HTML the server sends, not only injected by JavaScript after load.
26. Search Console's Pages report has been read this month and every "Not indexed" reason is either intended or fixed.

5. Content and trust

Explained in Chapters 3 and 6.

27. Each service has its own page; there is no single "Services" page trying to rank for everything.
28. Pages show first-hand experience: real numbers, real screenshots, real client situations (with permission).
29. Every article names a real author with a bio and a reason to trust them on the topic.
30. The About page, contact details, address and policies are easy to find from every page.

31. Old content is reviewed at least yearly; outdated pages are updated, merged or removed.
32. Nothing on the site was published just to hit a word count.

6. Local

Explained in Chapter 7.

33. Your Google Business Profile is verified, with the right primary category, hours, services and photos.
34. Name, address and phone (NAP) match exactly across your site, Google, Bing, Apple Maps and major directories.
35. You ask every happy customer for a review, and you reply to every review, good or bad.
36. Each location you serve from has its own page with genuinely local content, not a city name swapped into a template.

7. AI search (ChatGPT, Perplexity, Google AI Overviews)

Explained in Chapter 9.

37. Your robots.txt allows the AI crawlers you want citing you (OAI-SearchBot, PerplexityBot, Googlebot).
38. Key pages answer the core question in a clear, quotable paragraph near the top.
39. Facts that matter (prices, service areas, years in business, guarantees) are stated in text, not only in images.
40. Your business is described the same way everywhere it is mentioned: site, directories, social profiles, press.
41. You have checked what ChatGPT and Perplexity say about your business and your category, and noted what is wrong.

8. Structured data and measurement

Explained in Chapters 10 and 12.

- 42. Organization (or LocalBusiness) JSON-LD is on the site, with sameAs links to your real social profiles.
- 43. Structured data describes only what is visible on the page and passes Google's Rich Results Test; no self-written review stars.
- 44. You review Search Console clicks, impressions and top queries monthly, and you know which pages bring leads, not just traffic.

• • •

Free from davnoot.com/seo-checklist. Share it, print it, link to it.

Glossary

AI Mode. A conversational search experience in Google that answers questions with AI-written responses, running many related searches behind the scenes and linking to sources.

AI Overview. The AI-written summary that appears at the top of some Google results pages, with links to the pages it drew from.

Alt text. A short text description of an image, written in the page's HTML. Screen readers read it aloud, and search engines use it to understand the image.

Anchor text. The clickable words of a link. "Our emergency dentist page" is descriptive anchor text; "click here" is not.

Bing Webmaster Tools. Microsoft's free equivalent of Google Search Console, showing how Bing crawls and indexes your site. Bing's index also feeds Copilot and other AI services.

Canonical tag. A line in a page's HTML that names the main URL for that content, so search engines know which version to index when duplicates exist. Search engines treat it as a strong hint and can choose a different URL.

Citation (local). Any online mention of your business's name, address and phone number, such as a directory listing.

Core Web Vitals. Three measurements Google uses to judge page experience: LCP (Largest Contentful Paint, loading speed of the main content), INP (Interaction to Next Paint, responsiveness to taps and clicks) and CLS (Cumulative Layout Shift, visual stability).

Crawler. A program that visits web pages and follows links to discover content, a process called crawling. Also called a bot or spider. Examples include Googlebot, Bingbot, OAI-SearchBot and PerplexityBot.

Digital PR. Earning media coverage, and the links and mentions that come with it, by giving journalists useful expertise, data or stories.

Doorway page. A page created to rank for a narrow variation of a search, usually one of many near-identical pages such as city pages with only the city name changed. Google treats these as spam.

E-E-A-T. Experience, Expertise, Authoritativeness and Trustworthiness: the qualities Google's quality rater guidelines use to describe reliable content. It is not a single score, but it describes what Google's systems try to reward.

GEO and AEO. Generative engine optimization and answer engine optimization: names for the work of improving visibility in AI-generated answers. Most of it overlaps with good SEO.

Google Analytics 4 (GA4). Google's free analytics tool, which reports what visitors do on your site after they arrive.

Google Business Profile. Your free business listing on Google Search and Maps, which controls how you appear in local results and the map pack.

hreflang. HTML annotations that tell search engines which pages are language or regional versions of each other, so each searcher sees the right one.

Indexing. The process of a search engine analyzing a page and storing it in its index, the database it draws results from. A page that is not indexed cannot appear in results.

JSON-LD. The format Google recommends for adding structured data to a page, written as a block of code that does not change how the page looks.

Key event. In GA4, an action you have marked as important, such as a form submission or a phone tap. Previously called a conversion.

Keyword cannibalization. When two or more pages on the same site target the same search, competing with each other and weakening both.

Keyword map. A document assigning one primary keyword, and a few close variants, to each important page on your site.

llms.txt. A proposed plain-text file that gives AI tools a guide to a website's content. Some tools read it; Google has said it does not use it.

Manual action. A penalty applied to a site by a human reviewer at Google for breaking its spam policies, shown in Search Console under Manual actions.

Map pack. The block of (usually three) local businesses shown with a map near the top of Google results for local searches.

Meta description. A short summary of a page, written in its HTML, that search engines may show under the title in results. It affects clicks, not rankings.

NAP. Name, address and phone number. When your NAP is identical everywhere, search engines find it easier to trust your business details.

nofollow, sponsored, ugc. Labels added to links to tell search engines a link is not an editorial endorsement (nofollow), is paid (sponsored) or is user-generated (ugc).

noindex. An instruction in a page's HTML, or in its server response, that tells search engines not to show the page in results.

Redirect (301 and 302). An instruction that sends visitors and search engines from one URL to another. A 301 is permanent and transfers the old page's signals; a 302 is temporary.

Rich result. A search result with extra visual information, such as product prices, event dates or video thumbnails, made possible by structured data.

robots.txt. A file at the root of your website that tells crawlers which parts of the site they may visit. It controls crawling, not indexing.

sameAs. A structured data property that lists your official profiles on other sites, so search engines can connect them to your business.

Schema.org. The shared vocabulary used to write structured data, maintained jointly by the major search engines.

Search Console. Google's free tool showing how your site performs in Google Search: queries, clicks, impressions, indexing and errors.

Search intent. What a searcher is trying to accomplish: learn something (informational), compare options (commercial investigation), act (transactional) or find a specific place or brand (local or navigational).

Status code. The number a server returns with every URL request. 200 means OK, 301 means moved permanently, 404 means not found, 410 means gone on purpose, and 5xx means a server error.

Structured data. Code that describes a page's content and a business's details in a standardized format machines can read.

Title tag. The HTML element that sets a page's title, shown as the clickable headline in search results and in the browser tab.

XML sitemap. A file listing the URLs you want search engines to find and index.

About Davnoot

Davnoot is a digital marketing agency in Montreal. We work with small and mid-sized businesses, startups and online stores on SEO, AI search optimization, Google Ads, Meta Ads, email marketing and websites, in English and French.

Our approach is the one in this book. Fix the foundation first. Give every page one job. Do fewer things well, and measure them in leads and revenue rather than in impressions. We learned most of what is in here by doing it, for clients and on our own bilingual website, mistakes included.

We wrote this book because most of the businesses we meet have heard plenty of SEO advice, and very little of it came in plain language or in a sensible order. We hope it saves you time and money, whether or not you ever work with us.

The free checklist

The 44-item checklist from Appendix A is available as a free printable PDF at **davnoot.com/seo-checklist**. Print it, share it with your team, and link to it if it helps your readers. The full book is free as a PDF at **davnoot.com/seo-book**.

Talk to us

If you want a second pair of eyes on your site, you can book a free 30-minute call at **davnoot.com/book-call**.

davnoot.com