



GUIDE

FREE

HOW AI ASSISTANTS PICK THEIR SOURCES

Why some sites get cited and others never appear.
Observed patterns, labelled honestly.

01

AN HONEST FRAMING, FIRST

The providers publish little about how their systems choose sources. Everything in this guide is an observed pattern, drawn from provider documentation where it exists, from testing, and from watching what gets cited. Patterns, not rules, and they shift. Anyone selling you guaranteed AI citations is selling something the providers themselves do not offer.

With that said, the patterns are consistent enough to act on.

Assistants answer from two places, and the difference matters for you.

Training memory. What the model absorbed before its knowledge cutoff. You influence this slowly, over years, by being consistently present and consistently described across the web. There is no quick lever here.

Live retrieval. For current questions, the assistant searches, reads pages in the moment, and answers from what it found, often with citations. This is where you can act, and it behaves recognisably like search: get found, get read, get used.

Two hurdles follow. Your page has to be **retrieved**, reachable and relevant enough to surface. Then it has to be **cited**, useful enough in shape and substance that the answer draws on it. Most sites fail the first hurdle through crawl and rendering problems before content quality ever gets tested.

1 **Reachable and readable.** Sites blocking the AI crawlers in robots.txt, or serving key content only through JavaScript after page load, cannot be read in the moment. Retrieval favours what it can actually fetch. [Observed pattern]

2 **Answer-shaped content.** Pages that ask a clear question in a heading and answer it in the first sentence beneath give the system a clean, extractable unit. Buried answers extract badly. [Observed pattern]

3 **Specifics beat slogans.** Numbers, places, dates, named examples. Concrete claims appear in citations far more than positioning language does, which mirrors what a careful human quotes. [Observed pattern]

4 **Corroboration.** When third-party sites describe you consistently with how you describe yourself, the picture is confirmable. Assistants appear to lean on sources whose claims cross-check. [Observed pattern]

5 **Clear identity.** One business name, consistent details, a real about page, named authors on advice content. Ambiguous entities are risky to cite, and the systems behave accordingly. [Observed pattern]

6 **Freshness signals.** Visible dates, maintained content. For current-state questions, retrieval appears to prefer sources that show their age honestly. [Observed pattern]

WHAT DOES NOT APPEAR TO WORK, AND WHAT TO DO

Poor bets: keyword stuffing dressed up as "AI optimisation". Hidden text addressed to the models. Any service guaranteeing placement in AI answers. The old dark arts translate badly, since the selecting system reads more like a careful reader than a keyword index. [Observed pattern]

What to do instead: the companion checklist on the resources page turns these patterns into twenty-four concrete checks across access, structure, content, authority, schema and measurement. And start the ten-minute monthly habit today: ask the major assistants the five questions your buyers actually ask, log whether and how you appear, and date the log. Movement you cannot see is movement you cannot manage.



If you would rather have this done **with** you
than **to** you, this is the work I do.

START A CONVERSATION

PAUL PARSONS

CONSULTANT | DIGITAL MARKETING, AI & MARTECH

0403 127 478

www.paulparsons.com.au

pp@paulparsons.com.au

Perth, Western Australia 6111