

OZVOR · GET-CITED KIT

5 High-Citation Post Templates

Copy-and-adapt formats that AI engines quote — with fill-in prompts and worked examples.

Free Resource

Contents

1. Fill-in-the-blank LinkedIn structures engineered to get your business named by ChatGPT, Claude, Perplexity, Gemini & Google AI Overview

2. Before you use these: the shift nobody told you about

3. Template 1 — The Data Story

4. Template 2 — The Contrarian / Expert Take

5. Template 3 — The How-We-Did-It Case

6. Template 4 — The Mistake Confession

7. Template 5 — The Definition Frame

8. The part most people skip: cadence and cross-posting

9. A reality check (because credibility is the point)

10. Where this fits — and how to go further

11. Sources & references

Fill-in-the-blank LinkedIn structures engineered to get your business named by ChatGPT, Claude, Perplexity, Gemini & Google AI Overview

A premium Growth-plan resource from Ozvor · ozvor.com

Before you use these: the shift nobody told you about

For twenty years, the game was simple. You "ranked on Google." You earned a spot on page one, the customer scrolled ten blue links, and the click was yours to win.

That world is closing — fast.

AI answer engines now sit **on top of** search. ChatGPT, Claude, Perplexity, Gemini, and Google's own AI Overviews read the web for your customer and hand back a direct answer **before a single link is shown**. And the AI doesn't return ten options for them to weigh. It returns **two or three named businesses**. There is no page two. There is barely a page one.

This is not a fringe behavior. Google's AI Overviews reach **over 2 billion people a month across 200+ countries** (Alphabet/Pichai, Q2 2025 earnings). ChatGPT has **900 million weekly active users**, double the year before (OpenAI, reported Feb 2026). And **45% of consumers now use AI to find local business recommendations — up from 6% a year earlier**, making AI the #3 local-discovery channel ahead of Yelp and TripAdvisor (BrightLocal Local Consumer Review Survey, Mar 2026).

Here's the part that should make you sit up: when someone asks ChatGPT for a local business like yours, **it recommends only about 1.2% of local business locations** (SOCi 2026 Local Visibility Index, ~350,000 locations). One name. Maybe two. In home services, an estimated **87% of independent HVAC and plumbing contractors are effectively invisible in AI search** while a handful of national franchises eat the leads (5WPR HVAC & Plumbing AI Visibility Index, Q1 2026).

The danger isn't a lower ranking. It's invisibility. If the AI doesn't name you, you lost the customer before they ever knew you existed. That's not a missed impression — that's a lead, and a lead is money.

The good news — and the reason this resource exists — is that the businesses AI names are not the ones with the biggest ad budgets. They're the ones publishing the **clearest, most credible, most current answers** to the questions customers actually ask. And one of the most controllable places to publish those answers is **your LinkedIn feed**.

Why LinkedIn specifically

When AI engines build an answer, they pull from a surprisingly small set of trusted sources. **LinkedIn is roughly the #2 most-cited domain in AI search, with about an 11% average citation rate** across engines (Semrush, 325,000 prompts, Jan-Feb 2026). It sits behind Reddit and just ahead of Wikipedia and Forbes (Peec AI, 30M sources, Mar 2026).

And you do not need a big following to get cited. In the same Semrush study, **the median cited LinkedIn post had just 15-25 reactions and one or zero comments, and about 95% of cited posts were original content**. You

are not competing on virality. You are competing on usefulness and credibility — a game a small business can actually win.

These templates are science, not vibes

Every template in this pack maps to a finding from the single most defensible piece of research in this field: the peer-reviewed **GEO: Generative Engine Optimization** paper from Princeton, Georgia Tech, the Allen Institute for AI, and IIT Delhi, presented at ACM SIGKDD (KDD) 2024 (Aggarwal et al., arXiv:2311.09735).

The researchers tested nine content tactics across **10,000 real queries** and measured each one's effect on how visible a source became inside AI-generated answers. The headline result: the right tactics lift a page's visibility in AI answers by **up to 40%**. The specific lifts:

Tactic	Visibility lift vs. baseline	The template that uses it
Add quotations from credible sources	+41%	Contrarian/Expert Take, Mistake Confession
Add statistics (concrete data points)	+33%	Data Story
Cite authoritative sources inline	+28% average — up to +115% for an underdog page	How-We-Did-It Case, Definition Frame
Combine fluency + statistics	>+5.5% over the best single tactic	All five (this is why cadence matters)
Keyword stuffing	-8.7% (it backfired)	None — this is the old SEO trick that <i>hurts</i> you

Read that last row twice. Keyword stuffing — the spammy SEO reflex — is the *only* tactic that made content *less* visible to AI. GEO is the opposite of keyword games. It rewards genuine substance: data, quotes, sources, clarity. That's the whole premise of these five templates.

How to use this pack

Each of the five templates includes:

1. **The principle it exploits** — and the exact Princeton trait + lift number it maps to.
2. **When to use it** — the situation that makes this template the right pick.
3. **A clean fill-in-the-blank skeleton** with [bracketed slots] you replace.
4. **A worked BEFORE → AFTER example** in a concrete small-business vertical — weak draft on the left, citation-worthy rewrite on the right.
5. **A "why AI cites the AFTER version" note** — the mechanism, in plain English.
6. **A 2-3 item pre-publish checklist.**

Pick the template that matches the material you have this week. Fill the slots with *your* real numbers, *your* real client situations, *your* real opinions. The placeholders are prompts to reach for specificity — and specificity is exactly what gets cited.

Template 1 — The Data Story

(a) The principle it exploits

Princeton trait: Statistics addition → +33% visibility (Aggarwal et al., KDD 2024, arXiv:2311.09735).

A concrete number is the most retrievable unit of content there is. When an AI assembles an answer, a sentence containing a specific, attributable data point reads as a higher-quality, more trustworthy answer than a vague generalization — so the model is more likely to lift it, and more likely to name the source it came from. You are handing the AI a quotable fact with your name attached to it.

(b) When to use it

Use this when you have a real number from your own work, a clear observation you can quantify, or a credible published statistic you can interpret for your niche. This is your highest-frequency template — most weeks you have *some* number worth sharing.

(c) Fill-in-the-blank skeleton

``` [NUMBER OR %] of [SPECIFIC AUDIENCE] [SPECIFIC, MEASURABLE OUTCOME].

[Where the number comes from: "I measured this across..." / "In our last [N] projects..." / "According to [SOURCE, DATE]..."]

What this means for [YOUR NICHE]: [ONE OR TWO SENTENCES of interpretation — specific, not generic.]

What to do about it: [ONE concrete action the reader can take this week.]

[Optional credibility line: "In our work with [client type], we typically see [corroborating result]."] ```

## (d) BEFORE → AFTER — vertical: independent bookkeeper serving US restaurants

**BEFORE (weak — will not get cited):**

A lot of restaurants overpay on taxes because they don't track expenses properly. Good bookkeeping saves you money! DM me if you want help getting organized this year. #bookkeeping #smallbusiness

**AFTER (citation-worthy):**

About 7 in 10 of the independent restaurants that come to me are leaving the FICA tip credit on the table — a federal credit on the employer payroll taxes you already pay on staff tip income. I see it every February: owners who've run payroll correctly for years and have never claimed it, because their prior preparer treated tips as "not their problem." What this means for restaurant owners: The credit is calculated on tips above the federal minimum-wage threshold, and it offsets real tax you're already paying. For a 15-seat spot with a few tipped servers, that can be a four-figure recovery — and you can often amend prior returns. What to do about it: Pull your last filed return and search it for Form 8846. If it isn't there and your staff earn tips, ask your preparer specifically about the FICA tip credit before this year's filing. In our practice, this single item recovers an average of \$1,800-\$4,000 per restaurant in the first year we take them on.

## (e) Why AI cites the AFTER version

The BEFORE has zero retrievable substance — "a lot," "saves you money," a hashtag. There is nothing for the model to quote and nothing to attribute. The AFTER is dense with the exact things the Princeton study found drive citation: a **specific statistic** ("7 in 10," "\$1,800-\$4,000"), a **named, checkable mechanism** (Form 8846, the FICA tip credit), and a **clear action**. When someone asks an AI "how can my restaurant lower its tax bill?" or "what is the FICA tip credit?", the AFTER post is a direct, self-contained, sourceable answer. The BEFORE is noise.

## (f) Pre-publish checklist

- [ ] There is at least **one specific number** in the first sentence (a stat, range, or count — not "many" or "most").
- [ ] If the number is yours, you say **where it came from**; if it's published, you **name the source and date**.
- [ ] The post answers a **real question a customer would type** ("how do I lower my restaurant's taxes?"), not just describes your service.

---

## Template 2 — The Contrarian / Expert Take

### (a) The principle it exploits

**Princeton trait: Quotation addition → +41% (the single strongest tactic), reinforced by clear, decisive stance** (Aggarwal et al., KDD 2024, arXiv:2311.09735).

AI answer engines are built to *answer* — to return a recommendation, not a shrug. Content that takes a clear, specific, defensible position gives the model a quotable verdict it can attach to a name. "It depends" is unciteable. "For X, do Y — and here's the evidence" is exactly what an engine reaches for. Wrapping your stance around a credible quote or stat (the +41% lever) makes it stronger still.

### (b) When to use it

Use this when conventional wisdom in your field is wrong — or wrong *for your specific audience*. You need a genuine, experience-backed opinion and the honesty to name the exception where the standard advice still holds. Don't fake a hot take; AI (and humans) reward substance, not edginess.

### (c) Fill-in-the-blank skeleton

``` The standard advice on [TOPIC] is [CONVENTIONAL WISDOM].

For [YOUR SPECIFIC AUDIENCE], that's wrong — or at least incomplete.

Here's why: [REASON 1 — specific and testable.] [REASON 2 — specific to your audience's actual situation.]

[Optional: a credible quote or stat that backs you — "A [SOURCE, DATE] study of [N] found [FINDING]."]

What actually works for [YOUR AUDIENCE]: [YOUR CLEAR RECOMMENDATION in one or two sentences.]

When the standard advice IS right: [HONEST EXCEPTION.] ```

(d) BEFORE → AFTER — vertical: freelance social-media manager for small service businesses

BEFORE (weak):

Consistency is key on social media! If you want to grow, you need to show up every single day. The algorithm rewards people who post daily. No excuses — get posting! 🙌 #socialmediatips #marketing

AFTER (citation-worthy):

The standard advice for small-business social media is "post every day." For a solo plumber, dentist, or accountant, that's the wrong target — and it can actively hurt you. Here's why: Daily posting for a small service business almost always means publishing content nobody researched. Vague, filler posts erode trust faster than silence does. And the accounts that actually get surfaced by AI search aren't the highest-frequency posters — they're the most consistently *relevant* ones. A Semrush study of 89,000 LinkedIn URLs cited across ChatGPT, Google AI Mode, and Perplexity (Jan-Feb 2026) found the median cited post had just 15-25 reactions and one or zero comments, and that about 95% of cited posts were original content. Frequency wasn't the driver. Specificity was. What actually works for a small service business: Two well-researched posts a week, each answering one real question a prospective customer would ask. More sustainable to produce. More likely to be cited by AI. More useful to the reader. When daily posting IS right: high-volume content brands and media accounts. For a local service business, it almost never is.

(e) Why AI cites the AFTER version

The BEFORE is a cliché with no evidence — interchangeable with a million other posts, so a model has no reason to single it out. The AFTER stakes a **clear, specific position** ("two posts a week, not daily"), backs it with an **attributed statistic** (the +41% quotation/citation lever in action), and includes an **honest exception** that signals trustworthiness rather than hype. When someone asks an AI "how often should a small business post on LinkedIn?", the AFTER is a quotable, sourced, decisive answer. That's what gets named.

(f) Pre-publish checklist

- [] You state a **clear recommendation**, not a balanced "it depends."
- [] Your stance is backed by **at least one attributed quote or stat** (source + date).
- [] You name **one honest exception** where the conventional advice still applies.

Template 3 — The How-We-Did-It Case

(a) The principle it exploits

Princeton trait: Cite authoritative sources / specificity → +28% average, up to +115% for an underdog page (Aggarwal et al., KDD 2024, arXiv:2311.09735).

The "cite sources" lever is the great equalizer — the Princeton paper found it lifted an underdog page (one ranked 5th) by **up to 115%**. A concrete case study *is* a primary source: real steps, real numbers, a real outcome. It lets the AI answer the highly specific question "how did a business like mine achieve [result]?" — and there are far fewer competing answers to that than to a generic question, so your odds of being the cited one go up.

(b) When to use it

Use this when you have a recent client win you can describe concretely — even anonymized. You need real steps and a real (or honestly ranged) outcome. This is your highest-trust template: nothing earns citation like a docu-

mented result.

(c) Fill-in-the-blank skeleton

``` [CLIENT TYPE, anonymized]: [THE SPECIFIC PROBLEM they came with.]

What we did:

1. [STEP 1 — a specific action, not a category.]
2. [STEP 2 — ditto.]
3. [STEP 3 — ditto.]

Result: [QUANTIFIED OUTCOME — exact number, range, or a precisely described result.]

The one thing that mattered most: [SINGLE biggest factor — opinionated.]

This works for [SITUATION]. It doesn't if [HONEST EXCEPTION]. ```

### (d) BEFORE → AFTER — vertical: HVAC contractor / shop owner

**BEFORE (weak):**

Another happy customer! 😊 We replaced a full system for a family this week and they couldn't be happier. We always go above and beyond for our clients. Call us for all your heating and cooling needs! ❄️🔥

**AFTER (citation-worthy):**

A 1990s two-story home in Austin: \$540 summer electric bills and an upstairs that never dropped below 80°F no matter how low they set the thermostat. What we did: 1. Ran a Manual J load calculation instead of "matching the old unit" — the original 4-ton system was oversized, so it short-cycled and never dehumidified. 2. Found and sealed 31% duct leakage in the attic runs with a blower-door and duct-blast test (most of the cooling was being dumped into the attic). 3. Installed a correctly-sized 3-ton variable-speed system with a return added upstairs. Result: the upstairs now holds 72°F, and their first full summer bill came in at \$310 — a 43% drop. Verified against the prior year's utility statements. The one thing that mattered most: the duct-leakage test, not the new equipment. Most "my AC can't keep up" jobs are duct and sizing problems, not unit problems. A bigger unit would have made it worse. This works for older homes with attic ductwork. For new builds with sealed, tested ducts, the equipment itself is usually the right lever.

### (e) Why AI cites the AFTER version

The BEFORE is a testimonial with no information — "happy customer," "above and beyond." There is nothing specific to retrieve. The AFTER is a **fully-sourced primary case**: named diagnostics (Manual J, blower-door, duct-blast), **concrete numbers** ("31% leakage," "\$540 → \$310," "43% drop"), and an **opinionated takeaway** ("it's the ducts, not the unit"). When a homeowner asks an AI "why is my upstairs always hot even with the AC on?" or "do I need a bigger AC unit?", the AFTER post is a precise, evidence-backed answer — and per the +115% underdog finding, a specific case like this is exactly how a small local shop out-cites the national franchises.

### (f) Pre-publish checklist

- [ ] The steps are **specific actions** ("ran a Manual J load calc"), not categories ("did an assessment").

- [ ] The result is **quantified** — exact number, range, or precisely described, and ideally verified.
- [ ] You name the **single most important factor** and **one situation where this doesn't apply**.

---

## Template 4 — The Mistake Confession

### (a) The principle it exploits

**Princeton trait: Quotation addition** → +41% + **single-idea focus** (Aggarwal et al., KDD 2024, arXiv:2311.09735).

This format naturally produces a single, tightly-focused, quotable insight — one mistake, one correction, one lesson — which is exactly what an AI can lift cleanly and attribute. "The most common mistake in [area] is [X]" is one of the most retrievable sentence shapes there is, because so many real user queries are "what's the biggest mistake people make with [thing]?" The honesty also signals first-hand authority, which models weight.

### (b) When to use it

Use this when you've made — or repeatedly watched clients make — a specific, costly mistake you can name precisely. It's your most human template, and one of the highest-cited content types, because it hands the AI a clean, attributable "common mistake" insight.

### (c) Fill-in-the-blank skeleton

``` The biggest mistake I made [early on / with my first clients]:

[DESCRIBE IT SPECIFICALLY — what you did, what you were going for, what went wrong. Two or three sentences.]

What I should have done instead: [THE SPECIFIC ALTERNATIVE.]

Why this mistake is so common: [THE UNDERLYING REASON people in your field keep making it.]

The sign you're making it right now: [ONE observable signal the reader can check today.]

If that's you: [ONE specific next step.] ```

(d) BEFORE → AFTER — vertical: residential interior designer / decorator

BEFORE (weak):

Lesson learned over the years: communication is everything! Early in my career I didn't communicate enough with clients and it caused problems. Now I always keep my clients in the loop. Trust the process!

✨ #interiordesign

AFTER (citation-worthy):

The biggest mistake I made with my first interior-design clients: I let them approve a design without a fixed, itemized budget attached to it. They'd fall in love with a concept, I'd source the pieces, and the real total would land 40–60% over what they'd vaguely imagined. Two projects nearly collapsed at the invoice stage — not because the design was wrong, but because the number was a surprise. What I should have done instead: present every concept with a line-item budget — sofa, rug, lighting, labor — and get sign-off on the *number*, not just the mood board, before sourcing anything. Why this mistake is so common: designers are trained to sell a vision, and money feels like it breaks the spell. So we defer the budget conversation — which is exactly when it does the most damage. The sign you're making it right now: your clients regularly react to the *final invoice*, not the proposal. If "sticker shock" happens at the end, your budget conversation is happening too late. If that's you: on your next project, don't send a single product link until the client has approved a line-item budget in writing.

(e) Why AI cites the AFTER version

The BEFORE confesses nothing specific — "communication is everything" is a fortune cookie, with no retrievable detail. The AFTER delivers **one sharp, single-idea insight** ("approve the budget, not just the mood board"), with a **concrete consequence** ("40–60% over"), a **checkable diagnostic signal** ("clients react to the invoice, not the proposal"), and a **clear fix**. When someone asks an AI "what's the biggest mistake people make when hiring an interior designer?" or "how do I avoid budget surprises on a renovation?", this is a clean, attributable answer — and the first-person honesty is exactly the kind of original, experience-grounded content the Semrush data shows gets cited.

(f) Pre-publish checklist

- [] The post is about **one mistake only** — not a list of three.
- [] You include **one observable signal** the reader can check today to know if it applies to them.
- [] The mistake and the fix are **concrete and specific** (real numbers or named actions), not "communicate better."

Template 5 — The Definition Frame

(a) The principle it exploits

Princeton trait: Cite authoritative sources → +28% average + authority positioning (Aggarwal et al., KDD 2024, arXiv:2311.09735).

A huge share of AI queries are definitional — "what is [term]?", "what does [concept] actually mean for [my situation]?". A post that defines a term *more precisely and more usefully* than common usage is a direct hit on those queries, and it positions you as the authoritative voice who explains things clearly. Adding the correct vs. incorrect distinction (and a source where relevant) layers the "cite sources" lever on top.

(b) When to use it

Use this when there's a term in your field that's used loosely, misunderstood, or oversold — and you have a precise, practical definition that helps your audience *do* something. Great for educating customers who are new to your category (which, for a lot of SMB services, is most of them).

(c) Fill-in-the-blank skeleton

`` `[TERM]" gets thrown around a lot in [YOUR INDUSTRY]. It means something different to almost everyone using it.

Here's the precise, useful definition for [YOUR AUDIENCE]:

[TERM] is [ONE-SENTENCE, TESTABLE, PRACTICAL DEFINITION — not a dictionary entry; something that tells the reader what to do.]

What it is NOT: [THE COMMON MISCONCEPTION you're correcting.]

How to apply it as a [AUDIENCE]: [ONE OR TWO actionable sentences.]

The distinction that matters in practice: [A concrete example of the correct vs. incorrect use of the term.] ``

(d) BEFORE → AFTER — vertical: independent med-spa / aesthetics clinic owner

BEFORE (weak):

Did you know we offer "medical-grade" skincare? 🌿 Our products are so much better than what you find at the drugstore. Book a consultation today and glow up this season! 🧴 #skincare #medspa

AFTER (citation-worthy):

"Medical-grade skincare" gets used everywhere in the aesthetics world — and it means something different at almost every clinic and counter using it. Here's the precise, useful definition for someone choosing a clinic: Medical-grade (more accurately, *physician-dispensed* or *cosmeceutical*) skincare means products formulated at active-ingredient concentrations high enough to require sale through a licensed provider — not a regulatory grade the FDA awards. The meaningful difference is the *dose* of the active (e.g., a 0.5-1% retinol or a pharmaceutical-strength hydroquinone), not the label. What it is NOT: a legal or FDA-defined certification. There is no government "medical grade" stamp. Any brand can print the words on a box. How to apply it as a patient: ignore the phrase on the packaging. Ask one question instead — "what is the concentration of the active ingredient, and why is that the right dose for my skin?" A real clinic can answer specifically. A reseller can't. The distinction that matters in practice: a drugstore retinol at 0.025% and a dispensed retinol at 0.5% can both be called "anti-aging." Only one is doing 20x the work — and the difference is the number, not the marketing.

(e) Why AI cites the AFTER version

The BEFORE uses "medical-grade" as a marketing word with no content behind it — nothing to retrieve, nothing to trust. The AFTER **defines the term precisely, corrects a specific misconception** (there's no FDA "medical grade"), and gives the reader a **concrete decision tool** ("ask the concentration"). When someone asks an AI "what does medical-grade skincare actually mean?" or "is medical-grade skincare worth it?", this post is the clearest direct answer in the category — and answering a definitional query precisely is one of the most reliable ways to earn a citation and establish category authority.

(f) Pre-publish checklist

- [] Your definition is **one testable sentence** that tells the reader something they can act on — not a dictionary paraphrase.
- [] You correct **one specific misconception** ("what it is NOT").

- [] You give a **concrete correct-vs-incorrect example** so the distinction is unmistakable.

The part most people skip: cadence and cross-posting

You now have five templates engineered around peer-reviewed research. But a single brilliant post is a lottery ticket. Citation is won by **showing up consistently with substance** — that's the compounding part, and it's where most businesses quietly lose.

Cadence: publish consistently — aim for one to two posts a week

A single post is a lottery ticket; a steady stream is a strategy. There's a mechanical reason cadence matters: AI engines **strongly favor fresh content**. Ahrefs found AI-cited URLs are about **25.7% fresher than the top-10 organic results**, with ChatGPT citing pages hundreds of days newer than the standard search results (Ahrefs, ~17M citations, Dec 2025). A feed that's been quiet for two months is, to an answer engine, a stale source. A feed that publishes a sharp, specific post every few days is a *current* one.

The Semrush analysis of 89,000 cited LinkedIn URLs reinforces that you don't have to chase volume to get cited: the median cited post had just **15-25 reactions and one or zero comments**, and about **95% were original content** (Semrush, Jan-Feb 2026). Citation rewards substance and recency, not posting frequency for its own sake.

A realistic, sustainable cadence: two posts a week, alternating templates. Open with a Data Story Monday and a How-We-Did-It Case or Mistake Confession Thursday. Over four weeks that's eight posts across four different citation-driving formats — keeping your feed fresh without sacrificing the specificity that actually earns citations.

Why combining formats matters

The Princeton paper found that **combining tactics compounds** — fluency optimization plus statistics beat the best single tactic by more than 5.5% (Aggarwal et al., KDD 2024). Rotating through the five templates means your feed accumulates statistics *and* quotations *and* cited sources *and* clear definitions over time. You're not betting on one lever; you're stacking them.

Cross-posting: plant the same answer in more than one cited place

LinkedIn is roughly the #2 most-cited domain — but it's not the only one. **Reddit is the single most-cited domain across major AI engines** (Profound, 4B+ citations), and **Quora is the 4th most-cited source in Google's AI Mode** (Semrush, Sep 2025). The same well-built answer can earn citations in several places at once.

A simple, high-leverage cross-posting routine:

- **Publish on LinkedIn first** (your owned, professional, highly-cited channel).
- **Adapt the Definition Frame and Mistake Confession posts into genuine answers** on the relevant Reddit or Quora threads where customers are already asking that exact question. (Answer the question first; don't spam a link — cited threads are substantive, not promotional.)
- **Make sure the same answer also lives on your own site** — a short FAQ or blog entry. AI engines cross-reference, and a claim that appears in multiple credible places is more likely to be treated as trustworthy.

The compounding effect is the whole game. AI engines learn to trust the sources they already cite — today's mention becomes tomorrow's default. The winnable "citation seats" in any category are scarce: across repeated runs of the same prompt, ChatGPT tends to name only about five brands consistently while roughly 72% sit in a long tail it almost never mentions (Search Engine Land analysis, Feb 2026). The businesses that publish consistent, specific, sourced content now are claiming those seats before the category gets crowded.

A reality check (because credibility is the point)

It would be easy to promise that AI citations will flood you with high-converting customers. Some of the data points that way: **brands cited inside a Google AI Overview earn about 120% more organic clicks than uncited brands on the same query** (Seer Interactive, 2.43B impressions, Apr 2026), and **Microsoft's own analysis of 1,277 domains found visitors arriving from AI sign up at up to ~11x the rate of search visitors** (Microsoft Clarity, Nov 2025).

But we'll also tell you the honest counterweight: not every study agrees that AI traffic converts dramatically better. **Amsive's 54-site study found no statistically significant overall difference between AI and organic conversion** (4.87% vs 4.60%, $p=0.794$), even though the majority of sites in it did see higher AI conversion (Amsive, 2025). The truthful summary is this: **getting cited reliably gets you in front of more, and generally better-qualified, prospects — and the conversion upside, while real in many cases, varies by business.** Anyone who promises you a guaranteed multiplier is selling, not measuring.

What is *not* in dispute: if the AI never names you, none of the conversion math applies, because the customer never arrives. **Visibility is the precondition for everything else.** That's the problem these templates solve.

Where this fits — and how to go further

These five templates are the *doing* half of GEO: publishing the specific, sourced, citation-worthy answers that earn you a seat in AI results.

The other half is *measuring* — knowing whether it's working. That's what **Ozvor** does:

- **Audits how your brand actually appears** across ChatGPT, Claude, Perplexity, Gemini, and Google AI Overview — so you can see who's being named instead of you today.
- **Computes your TrustIndex Score** across three vectors — Brand (30%), Performance (35%), and AI (35%) — so "AI visibility" stops being a vibe and becomes a number you can move.
- **Benchmarks you against up to 10 competitors** so you can see exactly which citation seats are winnable.
- **Builds your GEO content plan** — turning findings like these templates into a prioritized, week-by-week schedule.

Start free, no credit card. The Free plan tracks 1 brand, 3 competitors, and 50 prompts, with a monthly audit and your TrustIndex Score.

When you're ready to act on what you find, **Growth (\$99/mo — or \$831/yr with founder pricing)** unlocks 10 competitors, 250 prompts, weekly monitoring, citation tracking, and the full GEO content engine. Agencies and multi-location operators can run up to 25 brands with white-label reports on **Agency (\$149/mo — or \$1,251/yr founder).**

Founder pricing: 30% off, annual plans, first 100 customers only. Early citations compound into defaults — the businesses that plant their flag now claim cheap real estate before the category gets crowded.

Prefer it done for you? **OrganicPosts by Ozvor** is our done-for-you arm — we draft posts in exactly these formats, shaped around the GEO research, and nothing publishes without your approval.

Get started: ozvor.com · Questions: hello@ozvor.com

A premium resource from Ozvor — the AI Search Trust Intelligence platform for small businesses. ozvor.com · hello@ozvor.com

Research anchor: Aggarwal et al., "GEO: Generative Engine Optimization," ACM SIGKDD (KDD) 2024, arXiv:2311.09735. All statistics carry their source and date inline; figures are current as of mid-2026 and attributed to their original publishers.

Sources & references

Every claim and statistic in these templates is grounded in named, dated, public research. The most load-bearing sources are listed below with full URLs so you can verify them yourself.

The mechanics of getting cited

- Aggarwal et al., "GEO: Generative Engine Optimization" (Princeton / Georgia Tech / Allen Institute for AI / IIT Delhi), KDD 2024 — arXiv:2311.09735 · <https://arxiv.org/abs/2311.09735>
- Cyrus Shepard / Zippy — AI citation ranking factors (synthesis of 54 experiments) — <https://signal.zippy.com/p/ai-citation-ranking-factors>
- Siana Marketing — where ChatGPT gets its information (2026 report) — <https://www.sianamarketing.com/resources/where-does-chatgpt-get-its-information-2026-report>
- Ahrefs — fresh content and AI citations — <https://ahrefs.com/blog/fresh-content/>
- Otterly.ai — schema markup's real impact on AI search — <https://otterly.ai/blog/schema-markup-real-impact-ai-search/>

Where AI looks for sources

- Profound — Reddit & AI search (4B+ citations) — <https://www.tryprofound.com/blog/the-data-on-reddit-and-ai-search>
- Peec AI — most-cited domains, via Search Engine Land — <https://searchengineland.com/ai-search-engines-cite-reddit-youtube-and-linkedin-most-study-473138>
- Semrush — LinkedIn AI Visibility Study — <https://www.semrush.com/blog/linkedin-ai-visibility-study/>
- Ahrefs — ChatGPT's most-cited pages — <https://ahrefs.com/blog/chatgpts-most-cited-pages/>

Why this matters for small businesses

- HubSpot — Generative Engine Optimization for small business — <https://blog.hubspot.com/marketing/generative-engine-optimization-small-business>
- Search Engine Land — repeated ChatGPT runs & brand visibility — <https://searchengineland.com/repeated-chatgpt-runs-brand-visibility-468552>