Get Paid to Fix AI Slop

By Joe Giler

Table of Contents

Chapter 1: What AI Slop Actually Looks Like: 12 Signatures You Can Spot in 30 Seconds

Let me tell you what this business actually is before we get into the craft of it, because if you understand the shape of the problem you'll never run out of work.

Between roughly 2023 and 2026, an enormous number of businesses published an enormous amount of machine-written text. Not because they were evil or lazy — because it was suddenly cheap. A blog post that used to cost $250 and take a freelancer four days now cost about eleven cents and took nine seconds. If you ran marketing at a mid-size company and someone showed you that math, you'd have done the same thing. Most of them did.

What nobody priced in was the cleanup. The text got published. It sat there. It ranked for a while, or it never ranked at all. It confused customers. It said things about the company's own product that weren't true. It gave legal advice on a law firm's site. It quoted a statistic that didn't exist and a journalist noticed. And now somebody inside that company has a spreadsheet with 400 URLs on it and a mandate from their boss that reads, roughly, "fix this."

That person is your customer. Not "businesses." Not "the market." One specific, mildly panicked marketing manager or ops lead who has a mess they didn't personally make and can't personally fix at volume. Your entire job is to be the person they hand the spreadsheet to.

To do that job you need one skill above all others: you need to be able to look at a page and know, fast, *what specifically is wrong with it*. Not "it feels like AI wrote it." That's useless. You can't invoice for a vibe. You need to name the defect, show the client the defect, and then charge to remove the defect.

This chapter gives you twelve named defects. I call them signatures. Each one is something you can point at on a screen with your actual finger. Each one is something a client can see once you show them, which means each one is billable.

Before the twelve: the honest caveat about "AI detection"

I need to be blunt about something here, because a lot of people selling courses in this space are not.

You cannot reliably prove a specific piece of text was generated by a machine. Not with GPTZero, not with Originality.ai, not with Copyleaks, not with any tool that exists as of 2026. These tools output a probability score, and those scores have two failure modes that matter enormously for your business:

- **False positives on real human writing.** Non-native English speakers get flagged constantly. So does clean, plain, well-edited prose — because "predictable word choice" is what these tools measure, and good editing makes writing more predictable, not less. There's a well-documented pattern of student essays being falsely flagged, which is why a number of universities disabled or de-emphasized these detectors in their integrity workflows.
- **False negatives on lightly edited machine text.** Run a generated draft through a paraphraser or change forty words by hand and most detectors drop below their own threshold.

So here is the rule I want you to internalize, and it will keep you out of trouble for your entire career in this niche:

Never sell "AI detection." Sell quality defects.

You are not a forensic examiner. You are not going to tell a client "this was written by ChatGPT" — because you can't prove it, because it might be false, and because if the writer was in fact a human freelancer, you've just accused someone of fraud on the basis of a probability score. That's how you get a furious email and lose the account.

Instead you say: *"This page has four unsupported claims, one broken internal link, a fabricated statistic attributed to a report that I can't locate, and a structure that doesn't answer the question in the title. Here's the markup. Here's what fixing it costs."*

Every single one of those is verifiable. None of them require you to guess at authorship. And here's the thing that makes this a real business rather than a gimmick: **the defects are the same whether a machine or a rushed intern**

produced them. You're selling editorial quality control. The AI boom just created an enormous amount of demand for it at once.

With that established — here are the twelve signatures.

Signature 1: The Hollow Opening

Read the first 100 words. Ask: *if I deleted all of this, would the reader lose anything?*

Machine-drafted articles almost always start by explaining that the topic exists and is important. "In today's fast-paced digital landscape, email marketing has become an essential tool for businesses of all sizes. As technology continues to evolve, understanding the fundamentals of email marketing is more crucial than ever."

Nothing happened in those two sentences. No fact, no claim, no stake, no specificity. It's throat-clearing rendered as prose.

How to spot it in 5 seconds: look for the pattern *[time-context phrase]* + *[topic]* + *[has become / is more important than ever]*. Common openers: "In today's," "In the ever-evolving world of," "As businesses navigate," "With the rise of," "In an era where."

What the fix is worth: this is one of the highest-value fixes you can make and one of the fastest. Cutting a hollow opener and replacing it with the actual first useful sentence measurably changes how long people stay on the page. You are moving the payoff from paragraph four to paragraph one.

The fix, concretely: find the first sentence in the article that contains real information — a number, a name, a step, a specific claim — and start there. Usually it's somewhere in the third or fourth paragraph. Everything above it goes in the bin. Then write one sentence of context in front of it if the reader genuinely needs orientation.

Before: "In today's competitive business environment, choosing the right CRM is a critical decision for companies of all sizes. With so many options available, it can be difficult to know where to start."

After: "Most small teams pick a CRM based on a demo, then discover eight months later that the thing they actually needed — a way to see every touchpoint with one customer on one screen — costs an extra $40 per seat. Here's how to check for that before you sign."

Same topic. One of them is a page. The other is throat-clearing.

Signature 2: The Symmetry Tell

This is the one nobody talks about, and once you see it you cannot unsee it.

Machine-generated structure is *too even*. Look at the article's outline. If there are seven H2 sections, and each one has exactly three bullet points, and each bullet point is roughly the same length, and each section runs 180–220 words — a human did not decide that. A human writing about seven things will have a lot to say about two of them, a moderate amount about three, and will barely mention the last two because they matter less.

Real expertise is lumpy. When someone actually knows a subject, their coverage is uneven in a way that reflects what matters. A plumber writing about pipe materials will spend 800 words on PEX vs copper and forty words on the rest, because that's the decision people actually get wrong.

How to spot it in 15 seconds: scroll the page fast and look at the shape of the sections in your peripheral vision. Uniform blocks = suspicious. Then check bullet counts. Three-three-three-three across every section is a machine's idea of "organized."

The fix: this one is the actual editorial work, and it's why you can charge for it. You have to decide what matters most and let the article become lopsided. Merge two thin sections. Blow one section out to triple length with real detail. Delete a section that only exists for symmetry.

When you propose this to a client, frame it as: *"Right now this page treats every subtopic as equally important, which means it doesn't answer the question the reader actually came for. I'll restructure so the highest-intent section gets the most depth."*

Signature 3: Unsourced Numbers

This is the one that gets companies sued, embarrassed, and de-indexed, and it's the single strongest argument for your invoice.

Machine-drafted text produces statistics that *sound exactly right*. "73% of consumers prefer brands that personalize their communications." "Companies using marketing automation see a 451% increase in qualified leads." "Studies show that 88% of users are less likely to return to a site after a bad experience."

Some of these numbers trace back to something real. Many trace to a decade-old vendor blog post that itself cited nothing. And some do not exist anywhere at all — they were assembled from the statistical shape of similar sentences.

How to spot it in 20 seconds: Ctrl+F the page for the percent sign. Every hit is a candidate. Then check three things:

1. Is there a source named at all? (Often: no.)
2. If a source is named, is it linked? (Often: no.)
3. If it's linked, does the link go to a page that actually contains that number? (Startlingly often: no.)

That third check is where you find gold. I've seen citations that link to a real, reputable organization's homepage — where the statistic appears nowhere. The link exists to create the *appearance* of sourcing.

The fix, in order of preference:

1. **Find the real source and link it precisely.** If the number is real, cite the actual study page, not a blog that mentions it.
2. **Replace with a number you can source.** Different stat, same point, real citation.
3. **Convert to an honest qualitative claim.** "Most buyers say personalization affects their decision" beats a fake 73%.
4. **Delete it.** If the sentence needs a fake number to work, the sentence doesn't work.

Why this is your best sales tool: unsourced numbers are the defect a client's boss can understand instantly. You don't have to explain editorial theory. You send them a screenshot with a red box around "451% increase in qualified leads" and the note "I could not locate any source for this figure; it appears on your pricing page." That email closes deals.

Signature 4: The Confidence Mismatch

Watch for text that is maximally confident about things that are genuinely uncertain, and maximally hedged about things that are simple facts.

You'll see: "The best CRM for small businesses is HubSpot." (Stated flatly. Depends entirely on the business.) And two paragraphs later: "It's generally recommended that you may want to consider potentially backing up your data." (Hedged into oblivion. Just back up your data.)

Humans with expertise calibrate. They say "almost always" and "in my experience" and "this depends on X" in the right places — and they say "do this" flatly when it's not actually debatable.

How to spot it: find every superlative — "best," "most effective," "the leading," "the top" — and ask "according to what, for whom?" Then find every stack of hedges — "may potentially," "it's generally considered that," "some experts suggest" — and ask "is this actually contested?"

The fix: flip the calibration. Superlatives get conditions attached ("HubSpot is the strongest option *if* you have a sales team of 5+ and need pipeline reporting; for a solo consultant it's overkill and Streak or a spreadsheet wins"). Over-hedged basics get stated plainly.

This is a fix clients feel immediately even if they can't name it. The page starts sounding like someone who's actually done the thing.

Signature 5: Definition Padding

The article promises advanced help and then spends 600 words defining terms the reader obviously already knows.

An article titled "Advanced Google Ads Bidding Strategies for Ecommerce" that opens with "What Is Google Ads? Google Ads is an online advertising platform developed by Google..." has a reader-model failure. Nobody searching for advanced bidding strategies needs to be told what Google Ads is.

How to spot it: compare the title's implied expertise level to the first two H2s. If the title says "advanced" and the H2 says "What Is X," you've found it. Also look for the "What Is / Why It Matters / Benefits Of" triple — three sections that collectively contain almost no instruction.

The fix: cut the definitions entirely, or compress them into one parenthetical, and if some readers genuinely need the primer, link to a dedicated beginner page. That last part matters — you've just created an internal linking opportunity and a reason to sell the client a second piece of content. Fixing pages profitably means noticing what else the site needs.

Signature 6: Zero First-Party Detail

This is the deepest defect and the one that most affects whether a page earns anything.

Read the whole article and ask: **could this have been written by someone who has never used the product, never talked to a customer, and never worked in this industry?** If yes, that's the problem. Not the phrasing. The emptiness.

Signs of zero first-party detail:

- No prices, or only vague ranges ("affordable pricing options available")
- No screenshots, or only stock photos of people pointing at laptops
- No specific version numbers, settings names, or menu paths
- No mention of what goes wrong or what the limitations are
- No dates, no "as of," no sense of when this was true
- Nothing that would be different if a competitor wrote the same article

The fix — and this one requires the client: you cannot manufacture first-party detail out of nothing. This is where you stop being a rewriter and become a content producer, and where your rate roughly doubles.

The move is a **20-minute expert extraction call**. You get someone at the company — the founder, a senior support rep, a salesperson — on a call and you ask:

- "What do customers get wrong about this, over and over?"
- "What's the number one reason someone churns off this feature?"
- "Walk me through the actual screens. What's the menu path?"
- "What does this genuinely not do well?"
- "Tell me about a specific customer situation last month."

Twenty minutes of that gives you more differentiated material than twenty hours of research. And the support rep is usually the best interview in the building — they know every failure mode by heart.

I'll cover the full extraction interview process later in this book. For now, just learn to see the absence.

Signature 7: The Fake Conclusion

Scroll to the bottom. If the last section is titled "Conclusion" and it restates the article's headings in paragraph form, you've found signature seven.

"In conclusion, choosing the right project management tool requires careful consideration of your team's size, budget, and workflow needs. By evaluating the options discussed above and considering your specific requirements, you can make an informed decision that will help your team stay organized and productive."

That paragraph has a job — filling space and hitting a word count — and it isn't helping the reader.

The fix: a real ending does one of three things. It tells you what to do next (a concrete first action, this week). It tells you what to watch out for (the failure mode most people hit). Or it makes a recommendation with a reason ("If you're under ten people, use Trello and stop researching. The tool isn't your bottleneck.").