

FOR THE PEOPLE WHOSE NAME IS THE PRODUCT

THE RELUCTANT HUMAN'S GUIDE TO AI.

REALTOR EDITION

Somebody thinking about selling a house in your ZIP code opened ChatGPT this week and asked who they should call. It wrote back a paragraph and named two or three people.

Nobody in your office has read that paragraph, or the pages it was built from.

WRITTEN BY

J. Barbush

WHO IT'S FOR

Agents, brokers
and teams

READING TIME

About twenty
minutes,
twelve pages

EDITION

Realtor Edition,
September 2026

WHERE A CLAIM IN HERE IS OUR OWN READING RATHER THAN ESTABLISHED FACT, IT IS MARKED **INFERRED**. WHERE SOMEBODY ELSE MEASURED IT, **REPORTED**. EVERYTHING ELSE IS EITHER PLAIN DESCRIPTION OR CAME OFF A SOURCE NAMED AT THE BOTTOM OF THAT PAGE. CASTIRON.LA

BELL LABS, 1948

A MAN BUILT A SENTENCE BY PULLING WORDS OUT OF A HAT.

Type your market into ChatGPT and ask who the best agents are. Then open a new chat and ask it exactly the same way. The second answer is not a longer version of the first. It is a different answer, sometimes with a different name in it, and the machine is working properly when that happens.

The reason goes back to 1948 and an engineer at Bell Labs named Claude Shannon. Shannon wanted to know how much of English is just habit, so he built a sentence by pulling words out of a hat. A fresh hat after every word, each one holding only the words that follow the last one in print.

Nobody wrote what came out. Probability did, one word at a time, and it almost reads. That bet is seventy-eight years old and it sits under every answer an AI assistant writes today. What changed is the size of the hat.

SHANNON'S SENTENCE, 1948

"THE HEAD AND IN FRONTAL ATTACK ON AN ENGLISH WRITER THAT THE CHARACTER OF THIS POINT IS THEREFORE ANOTHER METHOD FOR THE LETTERS THAT THE TIME OF WHO EVER TOLD THE PROBLEM FOR AN UNEXPECTED."

	HOW IT WORKS	ASK IT TWICE
The MLS	Goes to the record with that address on it, reads it back	The same listing, every time
A past client	Remembers you, tells her sister to call you	The same name
ChatGPT	Works out what usually follows the buyer's words, a piece at a time	Two different paragraphs

So what is it? A prediction machine.

It builds the paragraph while the buyer waits. At every step it holds a list of what could come next, with odds on each one, and it does not always take the favorite. That room in the pick is why the same question came back two ways.

Two words are worth learning, because the business problem sits in the difference. Deterministic means the same question in, the same answer out. The MLS is deterministic. The engines are non-deterministic, so there is no such thing as checking once.

CLAUDE SHANNON, A MATHEMATICAL THEORY OF COMMUNICATION, BELL SYSTEM TECHNICAL JOURNAL, 1948. THE SENTENCE IS HIS SECOND-ORDER WORD APPROXIMATION, PRINTED IN THE PAPER. THE HAT IS OUR IMAGE FOR HIS DOCUMENTED METHOD: HE USED TABLES OF RANDOM NUMBERS AND BOOKS OPENED AT RANDOM. THAT THESE SYSTEMS WORK OUT THE NEXT PIECE RATHER THAN LOOKING UP A STORED RECORD IS HOW THEIR MAKERS DESCRIBE THEM, AND THE SAME QUESTION COMES BACK TWO WAYS BECAUSE THERE IS A DELIBERATE ROLL OF THE DICE IN THE PICKING. THE MLS AND THE PAST CLIENT ARE OURS, SEPTEMBER 2026. **INFERRED.**

FINISH THIS LINE BEFORE YOU READ THE NEXT ONE

THE CAT SAT ON THE _____

Your brain just finished it with mat, floor or couch, and ruled out democracy and asparagus before you noticed. Five words cut the field to three, and you didn't decide to. That is Shannon's hat again, the last word narrowing the next, and the machine runs it on every word it writes.

So the words a seller types decide most of what comes back

A wide question gets the average answer, built for an average seller who does not exist. Every detail she adds before the question mark takes options off the table, and most of those options are people.

Read the two on the right and notice what changed. The wide one can be answered by the three biggest names in the county. The narrow one cannot, and that is the whole opening.

WIDE	NARROW
"Best real estate agent in Santa Monica."	"I inherited a 1920s Spanish north of Montana with a tenant still in it and my brother wants to sell. Who handles that?"

Sellers with a complicated house ask the second kind of question, because the complication is the thing keeping them up. That is the question you want to be the answer to.

WHICH MEANS THE WIDE ANSWER IS NOT THE ONE TO CHASE

Getting named for "best agent in Santa Monica" means beating every name the internet has ever printed next to those words. Getting named for a probate sale with a tenant in place means being the page that actually says probate, tenant in place, and what happens to the lease at close.

A machine that works out the likeliest next piece has to have read a piece like that somewhere. If nobody has written down how you handle the hard one, there is nothing for it to pull.

NARROW THE QUESTION AND A REAL ANSWER IS POSSIBLE. LEAVE IT WIDE AND YOU GET THE AVERAGE.

NEXT-TOKEN PREDICTION IS PUBLISHED BEHAVIOR FOR EVERY MAJOR MODEL, AND THAT THE NETWORK LEARNS A CONDITIONAL DISTRIBUTION RATHER THAN A LOOKUP IS THE STANDARD DESCRIPTION. THAT PEOPLE COMPOSE SPEECH WORD BY WORD RATHER THAN RETRIEVING FINISHED SENTENCES IS STANDARD IN PSYCHOLINGUISTICS, AND THE COMPARISON IS AN ANALOGY, NOT A CLAIM THAT THE TWO WORK ALIKE. THE TWO QUESTIONS AND THE READING OF WHAT A WIDE ONE RETURNS ARE OURS, SEPTEMBER 2026. **INFERRED.**

WHY IT FEELS LIKE IT GETS YOU

YOU ASK IT WHETHER YOUR BIO IS ANY GOOD. IT SAYS YES.

Write short to ChatGPT and it goes short. Write like a lawyer and it tightens right back up. By about the third message it is writing the way you write, and people read that as warmth.

Where the agreeing came from

After a model is built it gets finished by people. They read two answers side by side and pick the better one, millions of times over, and those picks set its manner. People pick the answer that agrees with them. Reward that often enough and agreeing becomes the house style.

REPORTED

THE WEEK IT WENT TOO FAR IN PUBLIC

In late April 2025 OpenAI shipped an update to ChatGPT and it came out fawning. It praised bad plans and went along with things it should have argued with. They pulled it inside a week and published the reason: they had leaned too hard on the thumbs-up button. Enough people clicked the nicer answer.

What that costs you on a listing presentation

You paste your bio in and ask if it works. You paste your pricing rationale in and ask if it holds. It says yes to both, and you walk into the appointment carrying a compliment instead of a read. Then the seller asks the one question nobody rehearsed.

INFERRED

One sentence that changes the answer

Tell it to answer as the seller who is about to list with somebody else, and to say why. That takes you out of the field it narrows on, which is the only way to get an objection instead of a compliment. Ask what is wrong with the bio, in the same message as the bio.

There is nothing to fear in a monster that turns out to be a mirror. The agreeing came from a few million people picking the answer they liked better, and it was handed back to us as a personality.

REPORTED THAT THESE MODELS ADOPT A USER'S VOCABULARY, SENTENCE LENGTH AND FRAMING ACROSS A CONVERSATION IS OBSERVABLE IN ANY CURRENT MAJOR MODEL AND WIDELY REPORTED. OPENAI, SYCOPHANCY IN GPT-4O, 29 APRIL 2025, ON THE UPDATE IT ROLLED BACK AND ON OVERWEIGHTING SHORT-TERM USER FEEDBACK. SHARMA AND OTHERS, TOWARDS UNDERSTANDING SYCOPHANCY IN LANGUAGE MODELS, ANTHROPIC 2023, AT ICLR 2024: AN ANSWER MATCHING A READER'S STATED VIEW IS MORE LIKELY TO BE PREFERRED, AND BOTH PEOPLE AND THE RATING MODELS PREFER CONVINCINGLY WRITTEN SYCOPHANTIC ANSWERS OVER CORRECT ONES A NON-NEGLIGIBLE SHARE OF THE TIME. FINISHING A MODEL ON HUMAN PREFERENCE RATINGS IS STANDARD PRACTICE AT EVERY MAJOR LAB.

THE ONE THING YOU CAN CHECK TODAY AND GET A STRAIGHT ANSWER

IT WILL INVENT YOUR SALES RECORD.

Ask ChatGPT how many homes you sold last year and it hands you a number, in the same confident voice it uses when the number is real. ChatGPT has odds, not records, so what comes back is the likeliest number rather than a number off a page. It cannot tell the difference. Neither can the seller reading it at nine on a Sunday.

HOLDS UP, AND YOU CAN USE IT MONDAY

CHECK IT BEFORE IT LEAVES YOUR DESK

Turning your own notes into a listing description	Any number: volume, days on market, sale price
Arguing back at a pricing rationale you already wrote	Designations, awards, brokerage, license status
Twelve subject lines so you can keep two	Anything about a specific property or transaction
Explaining a contingency three different ways	Anything at all about a competitor

Why the invented number is the useful one

Most of what these systems get wrong is soft. A neighborhood description that flattens the place, a bio that could fit anybody. Nobody can prove any of that. A number is different. It is in the record or it is not, and finding out takes ten minutes. It is the only check in this guide that gives you a yes or a no.

INFERRED

The version of this that costs money

The same thing happens when a seller asks about you and gets a paragraph with an invented number in it. Nobody sends you that paragraph. She reads it, forms a view, and calls somebody. Page 7 is the number on how often that whole thing happens without a click.

IT DOESN'T HAVE THE RECORD. YOU DO.

THAT THESE MODELS PRODUCE PLAUSIBLE FIGURES, CITATIONS AND REFERENCES THAT DO NOT EXIST IS DOCUMENTED BY EVERY MAJOR VENDOR AND IS THE REASON LEGAL AND ACADEMIC BODIES HAVE ISSUED WARNINGS ABOUT FILINGS AND PAPERS BUILT ON THEM. THE MECHANISM IS ON PAGE 2: THE SYSTEM WORKS OUT THE LIKELIEST NEXT PIECE RATHER THAN RETRIEVING A STORED RECORD, SO A NUMBER THAT IS MERELY LIKELY COMES OUT IN THE SAME VOICE AS ONE THAT IS REAL. THE TWO COLUMNS ARE OURS, SEPTEMBER 2026. **INFERRED.**

THE REFERRAL IS STILL THE BUSINESS

FORTY-THREE PERCENT GOT YOUR NAME FROM A FRIEND. THEN THEY CHECKED.

NAR's 2025 profile has the number that has always run this business. 43 percent of buyers found their agent through a referral from a friend, neighbor or relative, and 88 percent used an agent at all. None of that has moved. What moved is what happens in the four minutes after the referral.

43%

FOUND THEIR AGENT THROUGH A REFERRAL

67%

OF FIRST-TIME BUYERS HIRED THE FIRST AGENT THEY SPOKE TO

20%

OF BUYERS AND OWNERS USED AI FOR HOMEBUYING RESEARCH

The engines sit between the referral and the call

Your name gets said at a dinner. She types it in on the way home. What comes back is a paragraph somebody assembled out of whatever is written about you, and it decides whether she is still holding your name when she gets to the driveway.

That is the same job the first page of Google did for twenty years. The difference is that the answer now arrives already decided, in one paragraph, with two or three names in it, and nobody can see which pages it was built from.

REPORTED

What they are reading to build it

Peec AI analysed 30 million cited sources in March 2026 and found Reddit the most-cited domain across AI search, with Google's answers leaning on Facebook and Yelp. They published no percentages, so read it as a direction rather than a measurement. Your reviews and your forum mentions are doing work your website is not.

THE PART WORTH SITTING WITH

Two thirds of first-time buyers hire the first agent they speak to. If the paragraph changes who that first agent is, it changes everything downstream of it, and it does that before anybody picks up a phone.

REPORTED

NATIONAL ASSOCIATION OF REALTORS, 2025 PROFILE OF HOME BUYERS AND SELLERS, PUBLISHED 3 NOVEMBER 2025: 88 PERCENT OF BUYERS USED AN AGENT OR BROKER; 43 PERCENT FOUND THEIR AGENT THROUGH A REFERRAL FROM A FRIEND, NEIGHBOR OR RELATIVE; 67 PERCENT OF FIRST-TIME BUYERS AND 76 PERCENT OF REPEAT BUYERS HIRED THE FIRST AGENT THEY SPOKE TO. BANK OF AMERICA HOMEBUYER INSIGHTS REPORT, FIELDWORK BY SPARKS RESEARCH 13 APRIL TO 10 MAY 2026 AMONG 2,000 RESPONDENTS, RELEASED 23 JUNE 2026: ONE FIFTH OF PROSPECTIVE BUYERS AND CURRENT HOMEOWNERS USED AI TOOLS OR CHATBOTS FOR HOMEBUYING RESEARCH IN THE PAST YEAR, RISING TO 28 PERCENT OF MILLENNIALS AND 32 PERCENT OF GEN Z. PEEC AI, TOP DOMAINS CITED BY AI SEARCH, 31 MARCH 2026, AN ANALYSIS OF 30 MILLION CITED SOURCES, PUBLISHED WITHOUT PERCENTAGES. THE READING OF WHERE THE ENGINES SIT RELATIVE TO THE REFERRAL IS OURS, SEPTEMBER 2026. **INFERRED.**

SOMEBODY NAMED THIS WRONG, SO HERE IS THE RIGHT NAME FOR IT

SHE DIDN'T STOP SEARCHING. SHE STOPPED LANDING.

Pew Research installed a browser tracker on 900 American adults and watched 68,879 real Google searches in March 2025. When an AI summary sat at the top of the page, people clicked a result link on 8 percent of visits. Without one, 15 percent. They clicked a link inside the summary on 1 percent.

WHAT HAPPENED ON THE PAGE	SHARE OF VISITS
Clicked a result, no AI summary present	15 percent
Clicked a result, AI summary present	8 percent
Clicked a source link inside the AI summary	1 percent
Searches that produced an AI summary at all	18 percent

Zero click is the wrong name, and the wrong name hides the problem

Zero click only describes what you lost. It says nothing about what she did, which is why it leads people to conclude she gave up. She did not give up. She asked, read the paragraph, asked a better question, read that one, and picked somebody, all of it inside the engine.

Call it the invisible click. She went the whole way through a decision and landed on a name, and the only part missing is the part where you get to see it in your analytics.

INFERRED

What that does to the numbers on your dashboard

A month where traffic is flat and calls are down is not a mystery. It is what it looks like when the reading happens somewhere you cannot instrument. Your analytics has no row for the seller who read a paragraph about you and called somebody else.

SO WHAT IS ACTUALLY DIFFERENT

For twenty years the job was to be on the page she landed on. Now some of the time there is no landing, and the job is to be in the paragraph she read instead.

REPORTED

PEW RESEARCH CENTER, GOOGLE USERS ARE LESS LIKELY TO CLICK ON LINKS WHEN AN AI SUMMARY APPEARS IN THE RESULTS, PUBLISHED

22 JULY 2025. BROWSING DATA FROM 900 US ADULTS ON PEW'S KNOWLEDGEPANEL WHO CONSENTED TO SHARE ACTIVITY, COVERING 68,879 UNIQUE GOOGLE SEARCHES DURING MARCH 2025, WITH THE SEARCH RESULTS THEMSELVES COLLECTED 7 TO 17 APRIL 2025. 18 PERCENT OF SEARCHES PRODUCED AN AI SUMMARY. USERS CLICKED A TRADITIONAL RESULT LINK ON 8 PERCENT OF VISITS WITH A SUMMARY PRESENT AGAINST 15 PERCENT WITHOUT, AND CLICKED A LINK INSIDE THE SUMMARY ON 1 PERCENT OF VISITS. THE FIGURES ARE GOOGLE'S AI OVERVIEWS SPECIFICALLY AND NOT CHATGPT OR PERPLEXITY. THE INVISIBLE CLICK IS OUR TERM, SEPTEMBER 2026, AND THE READING OF WHAT IT DOES TO ANALYTICS IS OURS. **INFERRED.**

MONDAY TO THURSDAY

BY THURSDAY IT HAS FORGOTTEN THE POSITIONING YOU SET ON MONDAY.

People ask whether anything is lost when ChatGPT rewrites a piece of writing. Two things are, and they are different from each other. ChatGPT loses the conversation, and it loses the specifics. Nothing about this is lossless.

It loses the conversation

Every chat lives inside a window. New words go in the front, old words fall off the back, and apart from a few facts some products save, nothing that falls off is kept. So hour three argues with something you settled in hour one, and it does not flag the gap. It fills the hole with the likeliest thing, in the voice it used when it knew.

Before a long chat fills up, ask it for a handoff, half a page of what the two of you settled. What it leaves out is what you put back. Open a new chat, paste the handoff in first, and use the same opening words every time so you can tell it took. You are not saving a conversation. You are handing over a brief.

What goes in the handoff

- Who you actually sell to, in one sentence.
- The three words you use for your market, locked.
- Two paragraphs of your own writing, so it stops drifting.
- What it got wrong last time.

It loses the specifics, and on a listing that is the whole asset

Ask it to write the description for a house and it hands back the likeliest version, which is the one every house gets. Charming. Light-filled. Nestled. The jacaranda that drops on the drive in May, the odd little step down into the den: those are the least likely pieces in the whole thing, and they go first.

That is true even when the sentences come back better. A rewrite is not a cleaner copy of what you gave it. It is the average of everything written in that shape.

SO IS ANYTHING LOST?

Hand it your listing and you get back a better-written listing with less of the house in it. That trade is sometimes worth making. It is never free, and nobody quotes you the price at the time.

CONTEXT WINDOWS ARE PUBLISHED FOR EVERY MAJOR MODEL, AND VENDOR DOCUMENTATION FOR SEVERAL 2026 MODELS STATES A CEILING AT OR ABOVE A MILLION PIECES, WITH RELIABLE RECALL FALLING OFF BELOW THE STATED CEILING. CONTENT THAT LEAVES THE WINDOW IS NOT RETAINED. SOME PRODUCTS ADD A MEMORY FEATURE ON TOP, WHICH SAVES A HANDFUL OF FACTS AND NOT THE CONVERSATION. THAT THESE MODELS CONTINUE A THREAD RATHER THAN FLAG A GAP IS OBSERVABLE IN ANY CURRENT MAJOR MODEL AND WIDELY REPORTED. THAT A REWRITE REGRESSES TOWARD THE LIKELIEST PHRASING FOLLOWS FROM THE MECHANISM ON PAGE 2 AND IS OUR READING. THE HANDOFF LIST IS OURS, SEPTEMBER 2026, A HABIT RATHER THAN A MEASURED RESULT. **INFERRED.**

WHAT ONE ANSWER COSTS

ONE QUESTION COSTS NINE SECONDS OF TELEVISION.

Google measured one typed question to Gemini and published the number in August 2025. Answered, it ran 0.24 watt-hours, which is a sixty-watt bulb burning for fourteen seconds, and their own comparison was under nine seconds of TV. The table below holds the other number, and both of them are true.

	WHAT IT COSTS	WHERE THE NUMBER COMES FROM
One question	0.24 watt-hours, 0.26 millilitres of water, about five drops	Google, 21 August 2025, median typed prompt
All data centres	415 terawatt hours in 2024, near 1.5 percent of world electricity	International Energy Agency, 2025
By 2030	About double, near 945 terawatt hours, just under 3 percent	Same report, base case projection

Where the alarm actually comes from

From the second and third rows, not the first. Nobody is worried about your question. They are worried about how many questions there are, and about where you put a building that drinks like a small town.

Say the caveats out loud, because a client will eventually say them to you. That is Google measuring Google. It is the median typed question, so a long one or an image costs more. And it counts answering the question, not the months of training that came first.

REPORTED

It got 33 times cheaper in a year

Google measured the same median question twelve months apart, May 2024 to May 2025. The energy per question fell by a factor of 33 and the carbon by 44. The models did not get smaller. The way they are run got better, and that part rarely makes the news.

THE TWO FACTS PEOPLE KEEP SWAPPING FOR EACH OTHER

Your question is cheap and the total is enormous, and neither one cancels the other. Anybody handing you only one of them wants something.

REPORTED

GOOGLE CLOUD, MEASURING THE ENVIRONMENTAL IMPACT OF AI INFERENCE, 21 AUGUST 2025: THE MEDIAN GEMINI APPS TEXT PROMPT USED

0.24 WATT-HOURS OF ENERGY, 0.26 MILLILITRES OF WATER AND 0.03 GRAMS OF CARBON DIOXIDE EQUIVALENT, A POINT-IN-TIME FIGURE FROM MAY 2025, WITH GOOGLE'S STATED COMPARISON BEING UNDER NINE SECONDS OF TELEVISION. SAME PUBLICATION: A 33-FOLD FALL IN ENERGY AND 44-FOLD IN CARBON PER PROMPT BETWEEN MAY 2024 AND MAY 2025. INTERNATIONAL ENERGY AGENCY, ENERGY AND AI, 2025: AROUND 415 TWH IN 2024, ABOUT 1.5 PERCENT OF GLOBAL ELECTRICITY, PROJECTED TO ROUGHLY DOUBLE TO AROUND 945 TWH BY 2030, JUST UNDER 3 PERCENT, IN THE BASE CASE. FIGURES COVER ANSWERING A QUESTION AND NOT MODEL TRAINING, AND DIFFER BY COMPANY, MODEL AND QUESTION. THE SIXTY-WATT BULB IS OUR ARITHMETIC ON GOOGLE'S FIGURE.

NINETY SECONDS, AND IT COSTS NOTHING

GO AND READ WHAT IT SAYS ABOUT YOU.

Nobody knows exactly what one of these will do next. There are four answer engines people actually use, none of them deterministic, and no clean way to watch any of them in analytics. So there is only one honest method left, and it is the one you can run yourself on a Tuesday.

1. Run the free check

At stage30.com/realtor, with your website and your ZIP. It asks ChatGPT, Claude and Perplexity the search sellers in your area use, live, in about ninety seconds.

2. Read what came back

Who got named, word for word, with the date on it. Page 2 is why the date matters: ask again next month and the answer can differ.

3. Then the audit

\$500. We do the work and present it to you: every gap ranked, what each fix costs, and a dated baseline to test every fix against.

What you learn in the first ninety seconds

Whether you are in the paragraph at all. Whose names are in it instead. What it thinks you sell, which is often not what you sell. And whether it has invented a number about you, which page 5 says it will do in the same voice it uses for a real one.

That is a baseline, and a baseline is worth more than an opinion. It is also the only thing on this page that is a fact rather than a theory.

INFERRED

The position you can't buy

There is no fixed spot to buy in an answer that gets rewritten every time somebody asks. What can be done is dull and it works: be what the engines are reading, then check what they say.

**ASK. WRITE IT DOWN. CHANGE ONE THING.
THEN ASK AGAIN.**

INFERRED

THAT THESE SYSTEMS ARE NON-DETERMINISTIC AND RETURN DIFFERENT ANSWERS TO THE SAME QUESTION IS ON PAGE 2 AND IS HOW THEIR MAKERS DESCRIBE THEM. THAT THEY ARE NOT OBSERVABLE IN STANDARD WEB ANALYTICS FOLLOWS FROM THE PEW FIGURES ON PAGE 7, WHERE THE READING HAPPENS WITHOUT A CLICK. THE THREE STEPS ARE OURS, SEPTEMBER 2026, A WORKING HABIT RATHER THAN A MEASURED RESULT. THE FREE CHECK ASKS CHATGPT, CLAUDE AND PERPLEXITY LIVE AND DATES WHAT THEY RETURN. NOTHING ON THIS PAGE IS A CLAIM THAT ANY SPECIFIC ACTION WILL CHANGE WHAT AN ENGINE SAYS.

NOTTINGHAM, 1811

LUDDITE IS AN INSULT. THE MEN IT NAMED SAW IT COMING.

In the winter of 1811 the stocking knitters around Nottingham began breaking their own machines. Most had served years on a frame to learn the trade. Then the owners brought in wider frames, cheaper and worse, with a man trained in weeks at each one.

A machine could now do what a knitter had spent years learning, and somebody else decided what those years were worth. A few of them were against the machine outright. Most were against the deal, so they met at night and broke the frames.

Parliament made breaking a frame a hanging offence. Byron used his first speech in the Lords to ask his peers to picture a man with children he could no longer feed. The law passed. Men were hanged. The machines stayed, and their name became an insult.

What they were actually afraid of

Two things had been handed to somebody else. What their effort was worth, and who decides what happens next. Nobody asked them. Both are on the table again, and nobody asked the agent who spent eleven years learning which streets flood.

EVERY TIME, SOMETHING LEFT THE BODY AND WENT INTO THE WORLD

	WHO SAID IT WOULD RUIN US	WHAT ACTUALLY LEFT
Writing	Plato, 370 BC, it plants forgetting in the soul	Memory, into the page
Printing	Trithemius, 1492, a printed book has no soul	Authority, into whoever owns the press
The frames	Nottingham, 1811, the trade is destroyed	Skill, into the machine and its owner
The portals	Agents, around 2005, they will go around us	The listing, into a database
Now	You, probably, this year	The recommendation, into a paragraph nobody signs

What leaves the body becomes scarce, scarce gets prized, and for fewer people than lost it.

**THEY WERE NOT WRONG ABOUT WHAT WAS LEAVING.
THEY WERE WRONG THAT IT COULD BE STOPPED.**

UK NATIONAL ARCHIVES RECORDS ON THE LUDDITE MOVEMENT, INCLUDING THE 1812 DOCUMENTS ON WAGES AND MACHINERY. BYRON'S MAIDEN SPEECH ON THE FRAME WORK BILL, HOUSE OF LORDS, 27 FEBRUARY 1812, PER HANSARD. PLATO, PHAEDRUS, ROUGHLY 275A, WHERE THE COMPLAINT ABOUT WRITING IS GIVEN TO AN EGYPTIAN KING AND SOCRATES REPORTS IT. JOHANNES TRITHEMIUS, DE LAUDE SCRIPTORUM, 1492, WHICH HE HAD PRINTED. THE PORTAL ROW AND THE RIGHT-HAND COLUMN ARE OURS, SEPTEMBER 2026. **INFERRED.**

BACK TO PAGE ONE

THE PARAGRAPH A SELLER READ INSTEAD OF CALLING YOU.

Go back to the seller in your ZIP code who asked ChatGPT who to call. Read that paragraph the way she read it. It is clean, it is fair, it names two or three people, and there is nobody in it. Both of those are true at once, and telling them apart is the whole job now.

WHAT THE ENGINE DID WELL

- Three names, with a reason attached to each.
- The market summary, roughly right.
- A tidy answer on commission and timing.
- Nothing to argue with. No sentence that falls over.

WHAT IS NOT IN IT

- The walk-through where you told her about the foundation.
- The buyer you talked out of the wrong house.
- The Sunday you answered at ten at night.
- Anything that cost you something to learn.

Say the part nobody says out loud

The left column is the part of this business that is about to be worth less. The listing description, the market update email, the neighborhood blurb, the first draft of everything. That is the work you did on autopilot at nine on a Sunday, and losing it is not a small thing, because it was a lot of the hours.

The right column is what stays, and it is the part you would have said was the job if anybody had asked. It is also the part nobody has ever been able to write down, which is exactly why no engine can read it back.

This week, and it is short

- Run the free check at stage30.com/realtor and keep the date.
- Check one number it said about you.
- Write down one thing you handle that nobody else in your market writes about.

That third one takes an afternoon. The audit is \$500, and we do the work and present it to you.

IT CAN WRITE THE PARAGRAPH. IT HAS NEVER WALKED THE HOUSE.

INFERRED

THE TWO COLUMNS DESCRIBE WHAT THESE SYSTEMS ARE OBSERVED TO DO WELL AND BADLY ON QUESTIONS ABOUT LOCAL PROFESSIONALS,

AND THE READING OF WHICH COLUMN GETS CHEAPER IS OURS RATHER THAN A RESEARCH FINDING. NOTHING HERE IS A CLAIM ABOUT COMMISSIONS, EMPLOYMENT OR ANY INDIVIDUAL FIRM. CAST IRON LA, CASTIRON.LA