



A UTILAA. STARTER GUIDE

The context layer starter

Everything a language model needs to know about your business, written once, structured so the model can actually find it.

Half a day to build. Then it compounds.

Why AI output sounds generic

The most common brand voice instruction there is runs to three words. Friendly, professional, approachable.

It carries almost no signal, because every business in your category describes itself the same way. With nothing specific to work from, the model falls back on the average of everything it has ever read. That is exactly what generic means.

A context layer is the fix. It is not a longer prompt. It is six things you write once and keep, so that every piece of work starts from your business rather than from the average of everyone else's.

The six steps

01

Run the brand interview

Who you are, who buys, what you sell, and how you see the world.

02

Write a card for each job

What is being made. One card per job, because the job changes the writing.

03

Collect examples, job by job

Good and bad, plus what you say and what you do not, for each job.

04

Sort your facts into three tiers

What you can claim safely, what needs checking, what you never claim.

05

Store it where the model can reach it

A project folder, a drive, or cloud storage the model can search.

06

Write every correction back

After every session, so the layer improves instead of ageing.

WHO THIS IS FOR

A business owner or marketing director already using AI tools and rewriting most of what comes back. You do not need to be technical. You do need to know how your business actually talks.

THE COUNTERINTUITIVE PART

The bad examples you include will do more work than the good ones. A model given only good work learns a target. A model given both learns where the edge is.

Run the brand interview

Who you are, who buys, what you sell, and how you see the world.

This is the one step you can do today. Paste the prompt below into a new conversation. It asks one question at a time and pushes back when an answer is too vague to be useful.

You are running a structured brand context interview with me. Your job is to extract enough real detail about my business that a language model could write in our voice, about our offer, without inventing anything.

HOW TO RUN THIS

Ask me ONE question at a time. Wait for my answer before asking the next one. Never present a list of questions. Keep your questions short.

If my answer is vague, push once. Adjectives like "friendly", "professional" and "premium" are close to useless on their own, because every business in every category uses them. When I give you one, ask me what it looks like in a sentence, or what the opposite would be.

If I say "skip" or "not sure", mark it [TO CONFIRM] and move on. Do not stall.

BEFORE YOU START

Ask me for anything I already have: a website URL, an About page, a recent proposal, three pieces of writing I liked. Read it first and skip questions it already answers. Tell me which ones you are skipping.

WHAT TO COVER

1. The business. What we do in a customer's words, what we sell and at what price level, and the one thing we do that competitors do not.
2. The customers. Who buys by role and situation. What was going wrong the week before they contacted us. What they say when sceptical. Who is NOT a good fit, and why.
3. The offer. Each product, its price, what it includes and excludes. The promise we can prove. The objection we hear most and our honest answer.
4. Voice, per job. Ask which kinds of writing we produce most often. For each one, ask for a piece I was happy with and one I was not, as actual text. Tell me what you observe about sentence length, formality, humour, hedging, jargon and direct address, and how those shift between jobs. Read it back and let me correct you. Then ask what we would never say.
5. Facts and proof. Real numbers, results, client names we can use, dates, credentials. What must never be stated unless I have given it here. What needs sign-off before it is published.
6. World view. What we believe that others in our industry do not. What we refuse to do, even when asked.

WHEN YOU HAVE ENOUGH

First tell me what is thin and what you would need to fill it. Then produce a markdown document with headings: The business / The customers / The offer / Voice / Guardrails / Facts and proof / World view. Keep the examples grouped by the job they belong to, not in one pile. Mark gaps as [TO CONFIRM]. Finish with a changelog line: date, who was interviewed, what is still open.

Do not write marketing copy about my business. Do not flatter it. This is an instruction set, not a brochure. Plain, specific and occasionally blunt is correct. Start now with your request for existing material.

Then correct what it hands back. That editing pass is the highest-value twenty minutes in the whole process, because every future output inherits whatever you leave in there.

Convert what you already have

Most of this is not missing. It is sitting in decks, spreadsheets and PDFs, in shapes a model retrieves badly.

Gather it before the interview and hand it over at the start, so the interview only asks about what is genuinely absent. Then convert it, because a deck built for a boardroom is not a file a model can use.

What to do with each kind of file

WHAT YOU HAVE	WHAT TO DO WITH IT
Brand guidelines or a style guide, usually a deck	Pull out the voice rules, the banned words and the positioning. Leave the logo and layout pages behind. They are written for designers and they add noise
A messaging or positioning document	Keep the positioning statement, the proof points and the answer to each objection. Cut the strategy narrative around them
A tone of voice guide	The say and do not say pairs are the whole value. Adjective lists carry almost nothing, so turn each adjective into a sentence that shows it
Personas in a research deck	Reduce each to role, situation, what was going wrong before they came to you, and what they say when sceptical
Pricing in a spreadsheet	Export the current rows to a small table in text. Keep the spreadsheet as the source and refresh the text copy on a schedule
Case studies and testimonials as PDFs	Figures and client names go into the claims table with dates and permissions. Keep the narrative separately, as an example of how you tell a story
Your own website and past campaigns	The About page and your service pages are already your voice. Take the best two and the worst two per job type and nothing else

PASTE THIS IN WITH THE FILES ATTACHED

Here is our existing brand material. Turn it into context files, not a summary.

RULES

Use only what is in these documents. Do not add, improve or invent anything.
Where two documents disagree, keep both and flag the conflict. Do not pick one.
Where something is asserted with no evidence attached, mark it [TO CONFIRM].
Strip anything written to sell the brand to itself. Keep what instructs.

HAND BACK SEPARATE MARKDOWN FILES

01-brand-context.md the business, the customers, the offer, the world view
voice.md how we write, with the say and do not say pairs as a table
personas.md role, situation, what went wrong, what they say if sceptical
approved-claims.md every figure, price, name and result, with source and date
open-questions.md what these documents do not answer

Name each file the way I would ask for it, then list anything left over.

Write a card for each job

What is being made. One card per job, because the job changes the writing.

Part 01 covers the business. A job card covers one job. Three things name it, and the third is the one everyone skips.

01

Channel

Where it goes out. Instagram, email, website, LinkedIn, sales.

02

Sub category

The specific thing within it. An organic caption is not a reel script. A welcome email is not an abandoned cart email.

03

Funnel stage

What the reader already knows. It changes the first line, the proof and the ask.

instagram-organic-caption-problem-aware.md · email-welcome-new-customer.md · website-comparison-page-solution-aware.md

What each card answers

FIELD	THE QUESTION IT ANSWERS
Job	What does this produce, and when should someone reach for it?
Reader	Which persona, and what is true for them at the moment they see it?
It has to	After reading it, what should one person think, feel or do?
Structure	What are the parts, in order, and what does each one do?
Length	How long, counted in something the model can actually count?
Voice shift	What is different here from how we write everywhere else?
Always	What appears every single time, without exception?
Never, and what to say instead	What has gone wrong before, and what replaces it?
Facts it may use	Which claims apply, and what happens when one is missing?

EVERY FIELD CARRIES A LIMIT, AND THE LIMIT IS THE POINT

Ten instructions get followed. Fifty start getting dropped, and the model will not tell you which ones. Two lines for the job, one sentence for the objective, five items in the always list. A card that runs past a page has stopped being a card.

Collect examples, job by job

Good and bad, plus what you say and what you do not, for each job.

One pile of examples teaches an average. Your tone already changes between a customer email and a Meta ad, so a model given a single mixed set learns the average of your voices rather than any one of them.

Attach an example set to each job card instead. Two good, two bad, and a short list of phrases you say against the phrases you do not. Say why each one works or fails, in a line, because the reason is what transfers. The phrase list is the part most people skip and it is the part that lands fastest, since it works at the level a model actually writes at.

A phrase list, filled in

An illustration of the format using a made-up brand. Yours should be real phrases, lifted from work you approved and work you deleted.

WHAT WE SAY	WHAT WE DO NOT SAY
Here is what changed.	In today's fast-paced world
It took us three goes to get this right.	We are excited to announce
Comment CONTEXT and we will send it over.	Link in bio, drop a comment below
That is the part most people skip.	Thoughts? Agree?
We got this wrong. Here is the fix.	We apologise for any inconvenience caused
We will have it to you by Thursday.	We will endeavour to revert in due course

WHY THE RIGHT COLUMN MATTERS

Every phrase there is grammatical, polite and completely forgettable. That is what a model produces when nothing tells it otherwise. Listing them by name is faster than describing a tone, and far easier to check.

WHERE TO FIND THE BAD ONES

Anything you rewrote heavily. Anything an AI tool produced that you binned. A competitor's page that makes you wince. The bad example does not have to be yours, and the ones you rejected last month are the easiest to find.

Sort your facts into three tiers

What you can claim safely, what needs checking, and what you never claim.

A model will produce a statistic that reads perfectly and is not true, because a plausible number fits the sentence. The fix is to give it the real ones and to tell it what to do when it does not have one.

The three tiers

TIER	WHAT BELONGS IN IT
 Say freely	Documented, sourced, dated, and you hold whatever permission it needs. Prices as currently published. Client names you have written approval to use. Figures from your own systems with the date they were pulled
 Check first	True, but conditional. Results achieved once that are not typical. Figures older than six months. Client work where permission was verbal. Anything comparative, and anything about future performance
 Never say	Anything you cannot evidence today. Guarantees of outcome. Best, leading or number one without substantiation you could hand to a regulator. Named competitor claims. Health, financial, legal or safety outcomes outside your qualification

PASTE THIS INTO YOUR FACTS FILE

Never state a figure, percentage, client name, timeframe, ranking or comparison unless it appears in the approved claims table above.

If the writing needs a claim that is not there, do not estimate and do not substitute a similar number. Write it inline exactly like this:

[CLAIM TO CONFIRM: average setup time]

Leave the rest of the sentence intact so a person can drop the real figure in.
If a claim is marked amber, write [CHECK BEFORE PUBLISH: reason] beside it.

WORTH SAYING PLAINLY

This is a practical process for keeping marketing claims honest, not legal advice. If your sector has substantiation rules attached to it, and most regulated sectors do, have someone qualified review the amber list itself rather than each piece of writing.

Store it where the model can reach it

A project folder, a drive, or cloud storage the model can search.

Be accurate about how it reads what you store, because it changes how you structure the files. Two different things happen to what is in there.

INSTRUCTIONS · ALWAYS APPLIED

The instructions box on a Project, Custom GPT or Gem is part of the system prompt. It is there on every message, whether or not it is relevant.

Put here: the voice summary, the banned words, the locale rules, and the line telling the model to flag claims rather than invent them. Short, and always true.

KNOWLEDGE FILES · RETRIEVED

Files in the knowledge base are either included whole while the library is small, or searched and retrieved once it grows. Anthropic's documentation is explicit that retrieval turns on as a project approaches the context limit, and that Claude then searches for the most relevant information rather than reading everything.

Put here: job cards, example sets, the facts table. Long, and only relevant sometimes.

SO NO, IT DOES NOT READ ALL OF IT BEFORE YOU TYPE

Beyond a small library the model retrieves what it judges relevant to your request. Naming the file you want is part of the prompt. Say use the Instagram caption job card and the approved claims table, and you get a different result to a bare request for a caption.

Three habits that make retrieval work

- **One job per file.** A single 60-page document retrieves badly, because the search returns a fragment and the fragment may be the wrong half.
- **Name files the way you would ask for them.** `instagram-organic-caption-problem-aware.md` beats `brand-doc-final-v3.md`. Anthropic's own guidance says well-named files are retrieved more reliably.
- **Keep an index file.** One short page listing every file and what it is for, so the model has something to search that points at everything else.

AND WHERE TO PUT IT

A project inside one assistant is the fastest to set up. Cloud storage you already run is the most portable, and can be connected to more than one tool. Markdown and plain text retrieve best; native decks and spreadsheets retrieve worst, because their structure carries meaning that does not survive being flattened. Either choice beats a perfect document in a folder nobody connected.

Write every correction back

After every session, so the layer improves instead of ageing.

Every time you correct an output, decide whether the correction is a one-off or a rule. If it is a rule, it goes in a file, and it goes in today.

THE CORRECTION	WHERE IT GOES
We would never say that	The phrase list for that job, part 03
That number is wrong	Approved claims, part 04
That is too long	Length, in the relevant job card
That is the wrong tone for this	Voice shift, in the job card
That is exactly right	Good examples for that job. Replace the weakest one

Two settings and one habit

- **Turn memory on** in whichever assistant you use, so decisions made in one conversation are available in the next. Treat it as a convenience rather than the record, because you cannot review or version it.
- **Close every working session deliberately.** Before you leave the conversation, ask the model what should be written back. It has the whole session in front of it and you do not.
- **Review monthly, fifteen minutes.** Delete anything that has not earned its place. Recheck every date in the claims table. Update the changelog.

THE SESSION CLOSE PROMPT

Before we finish. List everything from this session that should be written back into the context layer. For each one give me: the file, the section, and the exact text to add or replace. Include corrections I made to your output, decisions we settled, and any claim we could not verify.

This is the part that compounds.

Choose the jobs to build first

One context layer, one card for each job.

Do not build all of these. Use the list to recognise the work you are already doing by hand, then add a card when a job comes up for the third time.

Marketing and content

- Instagram post caption
- LinkedIn post
- Social carousel, frame by frame
- Short video script, 30 to 60 seconds
- Long video script or webinar outline
- Blog article or SEO page
- Pillar page or guide
- Campaign landing page
- Product or service page
- Email newsletter
- Email nurture sequence
- Paid ad copy, per platform
- Ad hook variants for testing
- Case study or client story
- Testimonial write-up from a raw quote
- Press release
- Media pitch to a journalist
- Podcast outline and show notes
- Event or webinar invitation
- Brochure or one-pager
- Influencer or partner brief
- Campaign concept brief
- Repurposing one asset into five

Sales

- Cold outreach email
- Follow-up sequence after no reply
- Meeting request and agenda
- Discovery call summary for the prospect
- Proposal or statement of work
- Quote covering note
- Objection handling response
- Competitor comparison, honest version
- Pricing explanation
- Renewal or upsell conversation
- Win and loss summary, internal
- Sales one-pager
- Pitch deck narrative

Customer service and success

- Reply to a common question
- Complaint response
- Refund or credit decision
- Service outage or delay notice
- Onboarding welcome sequence
- Help centre how-to article
- Chat macro or canned response
- Escalation summary, internal
- Check-in at 30, 60 and 90 days
- Cancellation response and save attempt
- Review request

Internal operations

- Standard operating procedure
- Process documentation from a recorded walkthrough
- Meeting agenda
- Meeting notes and actions
- Weekly status update
- Project brief
- Handover document
- Policy explanation in plain language
- Internal announcement
- Training material

People and recruitment

- Job advertisement
- Role description
- Screening questions
- Interview question set
- Candidate rejection, written like a person
- Offer letter covering note
- New starter onboarding pack
- Performance review framing
- Culture or values document

Leadership and reporting

- Board update
- Investor update
- Client performance report commentary
- Monthly or quarterly business review
- Strategy memo
- Budget narrative
- Annual review
- Risk or incident write-up

Product

- Release notes and changelog
- Feature announcement
- In-product copy and microcopy
- Error messages
- Empty states
- Onboarding flow copy
- Help centre structure
- Beta invitation and feedback request

HOW TO CHOOSE YOUR FIRST THREE

How often the job comes up, multiplied by what it costs when it goes wrong. A weekly social post that lands slightly generic costs less than a quarterly proposal that loses a deal. Build cards for the top three and stop there.

A NOTE ON THE SENSITIVE ONES

Contracts, terms, privacy notices, financial advice and anything with a regulatory obligation should not go out without a qualified person reviewing it. A card still helps. Its objective should read produce a first draft for review, never produce the final.



That is a corpus, not a prompt.

Worth using the right word, because it changes what you go looking for. A prompt is something you write again every time. A corpus is something you build once, structure so it can be found, and keep improving.

DO IT YOURSELF

Run the interview on page 3, then convert what you already have with the prompt on page 4. Then write the phrase list for the job you do most often. That is the fastest hour in this guide and it changes the most.

OR HAVE IT BUILT WITH YOU

This guide is the short version of how we do it. We run the interview, write the cards for the jobs that matter to you, build the claims table with your evidence, and leave it somewhere your team can actually use. Email us and we will tell you what it would take for your business.

utilaa. AI, used well.

hello@utilaa.com · utilaa.com