AI FOR BUSINESS 2027

NOTES ON HOW I WORK.

The tools I use. The questions I ask.
The way I turn an idea into something real.

> I wanted you to see the process,
> not just the finished work.
> So I opened the notebook.

Goodness Business designer / Founder, Klarnow

I WANTED YOU TO SEE THE PROCESS.

I often begin with a voice note. An idea, a client conversation, something that annoyed me on a website, or three different things I need to connect. The first version is rarely tidy.

AI helps me work through that mess. I use it to pull context together, challenge the idea, write the brief, make the thing and check what happened next. But I still have to know what good looks like. I still have to make the decision.

This is the part I want to share with you: how I move from a rough thought to work somebody can use. I am opening the notebook: the tools, the inputs I give them, the prompts you can adapt and the checks I make before I trust an output.

You do not need my whole stack to begin. Start with one recurring piece of work. Finish it properly, learn from it, then add the next connection.

ABOUT THIS EDITION

This is my working stack as documented in November 2026, for taking into 2027. The prompts are teaching adaptations, not a transcript of every prompt I have used. This edition reflects the project work and records available for review.

OPEN THE NOTEBOOK.

Read it in order if you are setting up your way of working. If a job is waiting, go straight to that workflow and use the prompt.

CONTEXT + TOOLS / 04–07

What I give AI before the job. My actual working tools and the job I give each one.

THE THINKING / 08–09

CLIMB, Trinity and my 0–9 loop. Use the questions without making a small task complicated.

THE WORKFLOWS / 10–19

Briefs, meetings, proposals, design, images, websites, campaigns, research, recurring work and review.

PUT IT TO WORK / 20–25

A worked portfolio example, specialist roles, handoffs, the content desk and a seven-day start.

KEEP THIS NEARBY / 26–27

The working record, tool references and my last note.

TAKE WHAT YOU NEED

You can read the full guide without joining the waitlist. The editable prompt file is public too. Joining the waitlist is for programme and writing updates.

Download the copyable prompt notebook

GIVE IT THE BUSINESS BEFORE THE TASK.

If I ask for a website without explaining the business, I will get a website-shaped answer. That is why I keep coming back to context.

I want the AI to understand who I serve, what I am selling, what already exists and where the work is getting stuck. I also want it to know what I have already decided. Otherwise it can spend an entire conversation confidently reopening something I settled last week.

For each project, create one short working brief. Keep it current. I would rather have one accurate page than twenty pages of old instructions.

THE BUSINESS

Who I serve, what I sell, what is included and what the customer should do next.

THE REALITY

Current assets, people, budget or capacity constraints, evidence and working links.

THE DECISIONS

What is already settled, what has changed and what still needs a decision.

THE FINISH LINE

The deliverable, its owner, the deadline and the evidence that will count as finished.

MY ONE-PAGE CONTEXT SHEET.

I keep one current version for each project. Use the template below as a document, project note or the start of the conversation. Replace stale details instead of adding another layer of contradictory instructions.

TRY THIS / ADAPTED PROMPT

```
BUSINESS: [name, customer, offer, price, inclusions]
OUTCOME: [one result and target date]
CURRENT POSITION: [what exists, current links and files]
EVIDENCE: [dated results, feedback, relevant conversations]
PEOPLE + CAPACITY: [owners, time, budget ceiling]
DECIDED: [approved scope, brand rules, commitments]
OPEN: [uncertainties, missing inputs, dependencies]
DELIVERABLE: [what needs to exist next]
DONE MEANS: [observable acceptance evidence]
UPDATED: [date and what changed]

Read this before proposing work. Separate facts from assumptions. Identify
conflicts and missing evidence. Do not reopen settled decisions. Recommend the
next action for my outcome.
```

A NOTE FROM ME

Connected does not mean retrieved. Ask which emails, notes or files were actually read. If a source is unavailable, supply a relevant extract; do not let the AI invent the missing history.

WHAT IS ON MY DESK.

These are the tools I can point to in the work. They have different jobs. You do not need to buy all of them.

ChatGPT + voice input / my thinking workspace

I speak the rough version, bring the business context, then develop the brief, decision, copy or proposal. I correct the interpretation before asking for more output.

Granola / meeting memory

I return to the relevant note to recover decisions, proposals, responsibilities and unresolved questions. I check the source when exact wording matters.

Gmail + Google Workspace / original correspondence

Before drafting a reply, I bring in the actual thread. It anchors scope and commitments in what people said.

Working documents + Google Drive / continuity

Briefs, source assets, trackers and the latest approved version travel with the project. I ask for a specific file, not a vague claim that the AI remembers everything.

Connected apps / retrieval

Where a connection is available, I use it to bring the relevant material into the job. Access and permissions vary; a missing source remains a visible gap.

THE TOOLS AROUND THE WORK.

Codex / ChatGPT Work / ChatGPT Sites

I use these to move from a brief to a working site, document or system change. This portfolio was built and published through Sites.

ChatGPT image tools + Canva

Image generation and editing help me direct visuals. Canva is part of our team graphics workflow, including reusable event templates.

Figma / team collaboration

A design and handoff surface in the team. That is different from saying it is my personal tool for every daily task.

Git / GitHub + browser checks

Source history makes changes traceable. The browser shows whether the actual page and customer journey work.

Scheduled tasks / repeatable reviews

Recurring checklists and research reviews. A configured schedule is not evidence that its last run succeeded.

Behind relevant products / Resend, Stripe, OpenAI API

Email delivery, payment evidence and product AI. Klarnow Call also uses a telephony layer; outbound, WhatsApp and new integrations remain development work.

This is a working inventory, not a claim that every available app is part of my routine. I separate what I use, what the team uses and what we are still exploring.

HOW I MOVE THE WORK.

CLIMB is how I connect the business question to what we build and what we learn. I can move back a stage when the evidence changes.

C / CLARITY

What are we actually trying to achieve? Understand the business, purpose, customer and main uncertainty.

L / LINK

How do the audience, story, offer and customer journey connect? Find the handoff that is missing.

I / INITIATE

What is the plan? Choose the next outcome, owner, capacity and review date.

M / MAKE

What needs to exist? Build the assets and systems needed for the complete journey.

B / BOOST

What does the evidence say? Improve what works, change what does not and take the learning back to Clarity.

A NOTE FROM ME

An impressive output can still solve the wrong problem. I want a clear reason for making it before I ask AI to produce it.

TEST THE MOVE. THEN MAKE IT.

Trinity checks three things: **Commercial** – does it create or protect value? **System** – does it fit our people, tools and dependencies? **Execution** – can someone finish it, and how will we know?

TRY THIS / ADAPTED PROMPT

Review [decision] through Commercial, System and Execution. Use [current sources]. Identify the weakest assumption. Recommend proceed, modify, test, defer or stop. State the owner, next action, deadline and acceptance evidence. What new evidence would change the recommendation?

MY 0-9 LOOP

0 Unknown / new signal
1 See – observe the situation
2 Understand – diagnose it
3 Decide – choose the next move
4 Connect – bring the needed pieces together
5 Create – build the work
6 Operate – make the process usable
7 Reach – get it to the right people
8 Close – secure the next commitment
9 Prove – check the result, then learn again.

This is a process map, not a score out of ten. Clear Now covers 1–3; Build Now 4–6; Launch Now 7–9. Some jobs only need the relevant part.

VOICE NOTE TO A USEFUL BRIEF.

When I use this: I know something needs to happen, but my thoughts are coming out faster than the structure.

I speak freely first. Then I ask the AI to separate the outcome from the possible ways of getting there. That distinction matters: a new page, an event or a campaign is a proposed route. The goal might be qualified enquiries, repeat bookings or a stronger partnership.

01 Capture the rough thought without polishing it.

02 Add the relevant business context and recent decisions.

03 Ask for the goal, customer, constraint and next decision.

04 Correct the interpretation before asking for a full plan.

05 Turn the agreed direction into one deliverable with an owner.

TRY THIS / ADAPTED PROMPT

Here is my rough thinking: [voice-note transcript]. Turn it into a brief for [project]. Extract the outcome, intended audience, business value, constraints, decisions already made and unresolved questions. Separate my goals from my suggested solutions. Challenge anything unnecessarily complicated. Finish with the next concrete deliverable, its owner and what done means.

What I check: Did it understand the actual point? Has it invented certainty, dates or resources? Is the next action small enough to complete?

What I keep: The corrected brief, not every discarded version of the conversation.

A GOOD MEETING SHOULD CHANGE TOMORROW.

When I use this: We have had a useful conversation and I need it to change what happens tomorrow.

A tidy meeting summary is not enough for me. I want decisions, responsibilities and the things that would block the work. I also want to know whether something was agreed or merely suggested.

01 Open the relevant Granola note and the current project brief.

02 Identify what changed since the last conversation.

03 Separate decisions, proposals, questions and commitments.

04 Give each action one accountable owner and acceptance evidence.

05 Draft any follow-up against the actual email thread.

06 Put the next review in the working tracker.

TRY THIS / ADAPTED PROMPT

Read [meeting note] and [current brief]. Extract decisions, proposed ideas, commitments and open questions separately. For each action, show the owner, due date if explicitly agreed, dependency and evidence of completion. Mark anything unassigned or undated. Compare against the previous plan and show what changed. Draft a short follow-up without adding promises we did not make.

What I check: Who actually agreed to do it? Does the due date come from the conversation? Does "done" mean a file exists, a customer approved it, or something went live?

My rule: a proposed action stays proposed until someone owns it.

WRITE THE OFFER. CHECK THE COMMITMENT.

When I use this: I am preparing an offer, a partnership or a reply that affects scope and expectations.

I start with what the other person actually said. Then I connect it to the business problem, the value we can deliver and the work required. I do not want a polished proposal that quietly commits the team to work nobody priced.

01 Read the email trail and relevant meeting notes.

02 Identify the real need and the decision-maker's priorities.

03 Check the agreed scope, previous price and outstanding questions.

04 Define the deliverable, exclusions, dependencies and acceptance.

05 Draft the proposal in plain language.

06 Review the economics and commitments before sending.

TRY THIS / ADAPTED PROMPT

Using only [thread, meeting notes and approved offer], draft a concise proposal for [customer type]. Lead with their problem and the intended result. State deliverables, exclusions, responsibilities, timing, price or pricing questions, and acceptance criteria. Flag scope that has expanded beyond the existing agreement. Do not invent results, testimonials, discounts or client approval. End with one clear next decision.

What I check: Can we deliver it at the stated cost and capacity? Is the promise measurable? Has useful advice accidentally become unpaid implementation?

What I measure: Qualified conversations, proposals, agreements and collected revenue. Those are different stages.

SHOW IT WHAT YOU MEAN.

When I use this: I see a visual idea that could improve the work.

I do not just say "make it premium." I point to the part that matters: oversized type, spacing, a hover interaction, a restrained colour, how the image is cropped, or how the page asks for action. Then I connect that choice to my own brand and customer.

This portfolio started with an artist's page as a reference. I wanted the confidence and editorial rhythm, then adapted the content around business design, my work and the CLIMB framework. The reference gave us a direction. The business decided the message.

TRY THIS / ADAPTED PROMPT

Study this reference [image or URL]. Identify the specific layout, typography, spacing, imagery and interaction choices that create the effect. Apply the relevant principles to [my brand and goal], using my existing assets and approved copy. Explain what each choice helps the visitor understand or do. Preserve [brand rules]. Do not invent project results or substitute somebody else's identity for mine.

What I check: Does the page still feel like me? Is the offer understandable? Does the design work on a phone? Can the visitor see where to go next?

DIRECT THE IMAGE. CHECK THE DETAIL.

When I use this: I need a visual concept, a cleaner image or a specific correction.

My useful image prompts describe the change and the boundaries. If I want an interface icon removed from a screenshot, I say which icon. If the person, clothing, logo or event artwork should remain, I say that too.

01 Supply the real reference or edit target.

02 Name where the image will appear and the required crop.

03 Describe one clear change.

04 List the details that must remain intact.

05 Inspect the face, hands, lettering, logos and edges.

06 Review the image inside the actual page, not only by itself.

TRY THIS / ADAPTED PROMPT

Edit this image for [use]. Change only [precise element]. Preserve the person's identity, pose, clothing, existing brand marks, lighting and overall composition. Keep the text [exact text] unchanged. Remove [unwanted interface element] and reconstruct the underlying background naturally. Do not add new text or objects. Deliver a crop suitable for [placement].

What I check: A beautiful result can still be the wrong result. Compare it with the original. Reject changed logos, misspelled text and invented details.

BUILD THE JOURNEY. THEN TEST IT.

When I use this: The direction is clear enough to build or change a live experience.

I want the AI to inspect what already exists before it starts rewriting things. The job is to complete a customer journey. A hero section alone does not tell me whether the enquiry reaches anyone.

01 Give the approved brief, assets, existing site and constraints.

02 Ask it to map the current journey and identify the precise change.

03 Build the smallest complete version.

04 Test the actions, mobile layout, empty states and errors.

05 Review the actual page and make specific corrections.

06 Publish the reviewed version through the authorised route.

07 Check the live URL, domain, HTTPS and expected next step.

TRY THIS / ADAPTED PROMPT

Inspect [existing site or repository] before editing. Implement [bounded change] against this brief: [brief]. Reuse working parts and existing brand assets. Make the complete path from [entry] to [outcome] usable. Verify the affected behaviour, including mobile and failure states. Report what changed, what passed, what remains unverified and the actual published URL when deployment succeeds.

What I check: Does it save? Does refresh preserve it? Does the button do what the label promises? Is the email only drafted, submitted to a provider, or actually delivered? Can a visitor use the public page without my account?

A page that looks finished and a system that works are two different checkpoints. I need both.

CONTENT NEEDS SOMEWHERE TO GO.

When I use this: One event, offer or piece of work needs to become useful content across several channels.

I begin with the reason somebody should care and the action I want them to take. Then I adapt the material to the channel. I do not want the same paragraph copied everywhere with a different line break.

01 Identify the offer, audience, evidence and next action.

02 Choose the strongest source material: a project, conversation, clip or image.

03 Create a few distinct angles rather than endless variations.

04 Adapt each angle to LinkedIn, Instagram, TikTok or email as appropriate.

05 Check dates, venue details, links, names and rights to use the material.

06 Record the response and connect enquiries to follow-up.

TRY THIS / ADAPTED PROMPT

Turn [approved source material] into a campaign for [offer] aimed at [audience]. The next action is [CTA]. Give me one clear message, three useful angles and channel-specific drafts for [channels]. Use only verified details. Explain what evidence or asset each piece needs, who owns it and where replies should go. Keep the voice direct and recognisably mine. Do not add urgency or claims we cannot support.

What I check: Is this content making the offer clearer? Can people act on it? Who handles the response?

What I measure: Relevant replies, enquiries, registrations and paid outcomes where applicable. Views help diagnose reach; they do not settle the commercial question.

RESEARCH THAT EARNS ITS PLACE.

When I use this: I need to understand a product pattern, a market question or a current development before making a decision.

I use references actively. Sometimes I am studying a site's onboarding. Sometimes I am checking a partnership, a tool or a news story. I want the output to end in a decision or a useful change, not a pile of links.

TRY THIS / ADAPTED PROMPT

Research [specific question] for this decision: [decision]. Start with primary sources and show their dates. Separate verified facts, interpretation and unanswered questions. Compare only the dimensions that affect my choice: [criteria]. Explain what is transferable to our context and what depends on assumptions we have not tested. Finish with the smallest next experiment and the evidence that would confirm or reject it.

What I check: Are the sources current? Is a competitor's marketing claim being repeated as fact? Are we copying a feature when we should be solving the underlying customer problem?

My note: research should reduce the uncertainty that blocks action.

IF I REPEAT IT, I WRITE THE PROCESS.

When I use this: I notice myself requesting the same review every day or week.

I use recurring checklists and research tasks to bring the next decision back into view. I do not treat the existence of a schedule as proof that something happened. A good recurring task has a clear source, a bounded job and an output I can inspect.

01 Perform the task manually enough to understand it.

02 Define the trigger, sources and exact output.

03 State what it can update and what still needs my decision.

04 Add a clear rule for missing context or failed access.

05 Review the result and adjust the instructions.

TRY THIS / ADAPTED PROMPT

At [time and timezone], review [authorised sources] for [purpose]. Update [working checklist or report]. Mark completed work only when there is evidence. Carry open work forward without duplicating it. Highlight changed commitments, blockers and the three highest-impact next actions. If a source is unavailable, say so and keep the uncertainty visible. Do not send messages or make external commitments unless separately authorised.

What I measure: Missed follow-ups, time spent reviewing, actions completed and whether the output still helps me decide. If the report becomes noise, I change or stop it.

BEFORE I TRUST THE OUTPUT.

I can ask the AI to review its work, but that does not make the review independent proof. I still need the source, the real page, the actual record or the result of a test.

My check is simple:

01 **Truth**: Which claims come from evidence? Which are assumptions?

02 **Fit**: Does this solve the problem I actually gave it?

03 **Voice**: Does it sound like me and fit the brand?

04 **Function**: Does the promised action work end to end?

05 **Ownership**: Who handles the next step, including the failure?

06 **Value**: Was this worth the time and cost?

TRY THIS / ADAPTED PROMPT

Review this deliverable against [brief and acceptance criteria]. Identify factual errors, unsupported claims, broken journeys, unclear ownership and unnecessary complexity. Rank only material issues. For each, show the evidence, the practical consequence and the smallest fix. Apply authorised fixes, verify them and report remaining limits. Do not call it complete because the draft exists.

A useful distinction: drafted, approved, published, received and paid are different states. I want the state the evidence supports.

THIS PAGE STARTED AS A ROUGH THOUGHT.

Here is a real example you can inspect: my personal portfolio. The starting point was a spoken brief for a premium personal page that positioned me as a business designer, showed the work and invited people into AI education.

01 **Clarity**. I wanted the page to connect my work, thinking and offer. We used the available history and project evidence to avoid positioning me from visual clues alone.

02 **Link**. I supplied an artist-page reference. We translated its strong typography and editorial layout into my name, portraits, work and CLIMB.

03 **Initiate**. We chose the page sections and a clear next action: the AI guide and programme waitlist.

04 **Make**. We built the page, project details, writing section, guide and signup. I reviewed the images and corrected specific parts.

05 **Boost**. We published the site, connected the domain and checked the live page. Conversion results still need real visitor data; publishing alone is not proof of sales.

A NOTE FROM ME

The later request was precise: DAB gets the supplied campaign image; SummerYard keeps its artwork but loses the Instagram controls; Klarnow gets its brand image. Clear feedback makes revision faster.

GIVE EACH SPECIALIST A REAL JOB.

For a larger job, I ask for distinct perspectives or bounded specialist work. That can help with research, copy, assets or review. I still want one person or process accountable for the combined result.

I do not need several agents for a simple correction. I use them when the work can be separated cleanly and the findings can be checked against the same brief.

TRY THIS / ADAPTED PROMPT

Use this approved brief: [brief]. Separate the work into [research / copy / assets / implementation / review] only where useful. Give each role a bounded deliverable, sources, constraints and definition of done. Name one owner to reconcile disagreements and integrate the outputs. Do not let multiple workers overwrite the same files. Report assumptions and unresolved conflicts before calling the combined result finished.

THE HANDOFF I WANT

What changed. Where the source lives. Who owns the next step. What passed. What remains unverified. Which decision is needed. If those are missing, the handoff is still incomplete.

A NOTE FROM ME

Asking for a "25-year expert" can describe the standard I want. It does not create real experience or make an answer correct. The work still has to earn my trust.

ONE SOURCE. SEVERAL USEFUL PIECES.

My news and content work starts with a current source or something we have actually made. I want the research, article, social copy and next action to connect.

01 **Capture.** Save the primary link, date, source material and the question our audience is asking.

02 **Decide.** Identify why it matters to a business owner and what we can say from evidence. Skip a story if it does not earn its place.

03 **Write.** Build the full article or note first. Then pull a useful argument for LinkedIn, a visual sequence for Instagram and a spoken angle for a clip.

04 **Package.** Keep the headline, approved copy, image, source links, channel, owner and destination together in a content record.

05 **Review.** Check facts, brand, links and the actual published status. Track useful responses and follow-up, not just the number of drafts.

TRY THIS / ADAPTED PROMPT

Using [verified source], produce a content record: audience question, main point, source and date, article draft, LinkedIn draft, Instagram concept, short-video outline, asset needed, CTA, owner and review status. Clearly label draft versus approved versus published. Do not manufacture facts to fill a content calendar.

START WITH ONE FINISHED LOOP.

If I were starting this way of working from scratch, I would begin with a chat workspace, one current business brief and one recurring job. Add meeting notes when meetings are the source. Add a coding environment when there is a real build to do.

Before adding a tool

What job will it do? What current work does it replace? Who will use the result?

Before paying for another subscription

Can the existing stack do the job? Will the benefit exceed the subscription, setup and review time?

Before automating

Can I describe the inputs, trigger, output, exception and owner? Does the manual version work?

Before publishing

Is the claim supported? Is this the approved version? Do I have the right to use the material? Does the next action work?

A NOTE FROM ME

I can share the method without sharing somebody else's private inbox. Use anonymised examples, remove credentials and give a tool only the access its job needs. Specific workflows are useful; private customer data is not part of the lesson.

MAKE YOUR FIRST THREE DAYS COUNT.

Pick a real job that is already waiting. A client update, an enquiry response, an offer page or a meeting follow-up is enough. Write down the current result and time spent so you have something to compare.

DAY 1 / GIVE IT CONTEXT

Fill in the context sheet. Add current sources and remove outdated decisions. Ask the AI what is confirmed, missing or contradictory.

DAY 2 / PICK THE BOTTLENECK

Use CLIMB and the Trinity questions. Choose one outcome, the relevant workflow and the evidence you will look for.

DAY 3 / RUN THE WORKFLOW

Use a prompt from this notebook with your own inputs. Keep the first output and record what you corrected. Finish the chosen deliverable before expanding the plan.

TRY THIS / ADAPTED PROMPT

My chosen job is [job]. My current baseline is [time, quality or outcome]. Using this context [brief], propose one small test for the next seven days. Name the owner, inputs, output, checks and review date. Keep the scope within [capacity].

FINISH. USE IT. LEARN SOMETHING.

DAY 4 / CHECK IT

Use the review prompt. Verify the claims, inspect the real output and test the whole journey. Correct the material issues.

DAY 5 / PUT IT TO WORK

Use the reviewed deliverable in the appropriate channel. Someone must own the response, delivery and failure cases.

DAY 6 / SAVE THE RECIPE

Keep the prompt, context requirements, source location and acceptance checks. Add a schedule only if the task is stable and worth repeating.

DAY 7 / REVIEW THE RESULT

Compare with the starting point. Record time spent, errors, useful responses and commercial outcomes where relevant. Decide keep, change or stop.

A NOTE FROM ME

A proposal is not a sale. An invoice is not cash collected. A scheduled post is not a published post. Write down the state that the evidence actually supports.

If the result is weak, go back to the input and the problem before buying another tool. The most useful change may be a clearer offer or a better handoff.

THE WORKING RECORD.

Use the same record for each run. It makes the learning visible, and it gives the next person enough context to continue.

THE JOB + OWNER

What needs to happen, for whom, by when?

THE INPUT + TOOL

Which current sources and prompt did I use? Which tool did which part?

THE OUTPUT + CORRECTIONS

What was produced? What did I have to fix, and why?

THE PROOF

What actually happened? Keep the approved file, live link, test result or provider record.

THE COST + RESULT

Time spent, review time, direct cost and useful outcome. Use consistent definitions.

THE NEXT RUN

Keep, change or stop? Update the source, prompt or process. Set the next review date.

Example calculation: if a comparable job took 90 minutes before and now takes 35 minutes plus 20 minutes of review, that is 35 minutes saved. This is an illustration, not a claim about my results.

A LAST NOTE FROM ME.

I am still refining this way of working. The tools change. The responsibility does not.

Bring the context. Ask a better question. Make something useful. Check it. Put it in front of real people. Learn from what they do.

That is the part I want you to take with you.

GOODNESS

TOOL REFERENCES / CHECKED NOVEMBER 2026

ChatGPT Work and Codex

Projects in ChatGPT

Connected apps in ChatGPT

Scheduled tasks in ChatGPT

Granola: getting started

Canva: editing images with Magic Edit

These links explain the products; the workflow descriptions come from my working practice. Features and access can change by plan, account and date. This guide does not predict 2027 product features.

Written with AI assistance from my project context, available records and editorial direction. The teaching prompts are adaptations. Private correspondence and client data are not reproduced here.