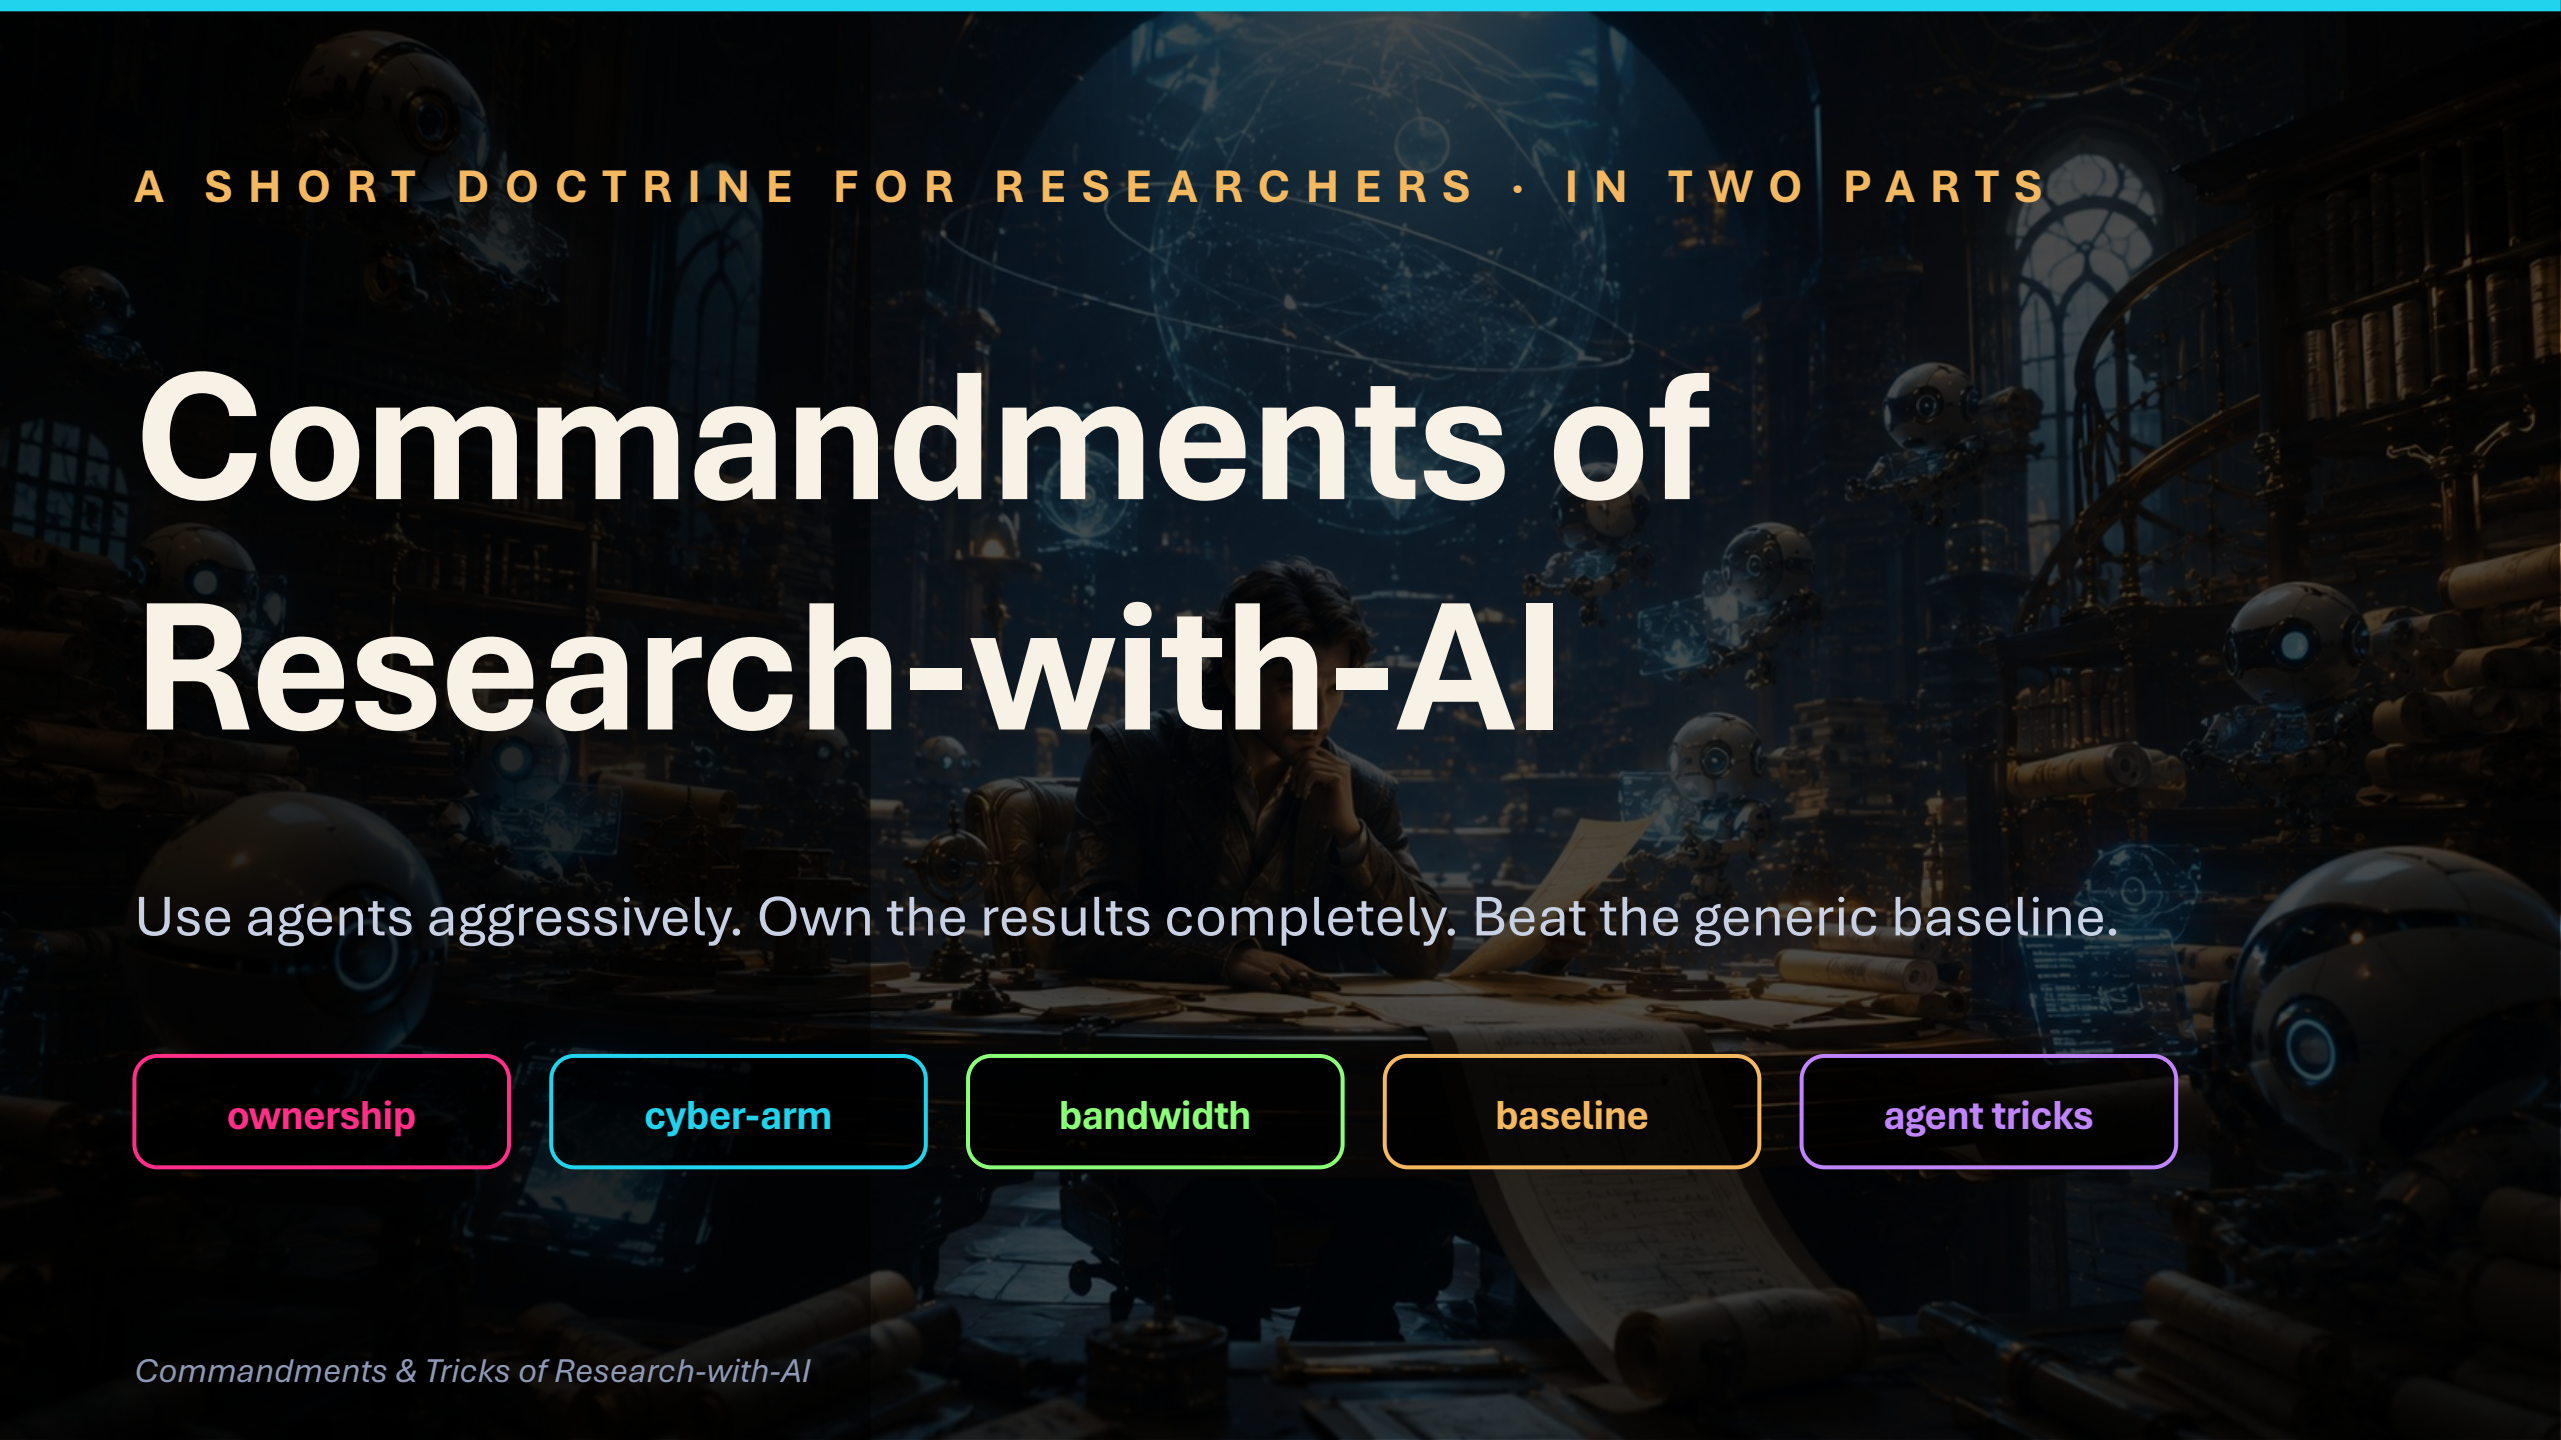

A SHORT DOCTRINE FOR RESEARCHERS · IN TWO PARTS

Commandments of Research-with-AI

Use agents aggressively. Own the results completely. Beat the generic baseline.

ownership | cyber-arm | bandwidth | baseline | agent tricks

PART I

The commandments.

The stance: four rules for how a researcher stands relative to AI.

I **AI is our minion.** — ownership

Own every claim, or do not use it.

II **AI is not your peer. It is your cyber-arm.** — cyber-arm

Attach it to the work you care about most.

III **Our input–output bandwidth is the bottleneck.** — bandwidth

Specification in. Comprehension out.

IV **Be above AI on everything you think about.** — baseline

Generic AI output is the floor, not the ceiling.

Part I sets the stance. Part II makes it run.

COMMANDMENT I

AI is our minion.

Own every claim, or do not use it.

An agent output is not research. It becomes research only when a human can understand it, defend it, and communicate it.

PRACTICAL STANCE

If you cannot explain it, it is vapor.
If you doubt it, make the agent justify it.
If it is too hard, demand a simpler form.

Demand adaptation.

Adapt the minion to your cognition — not the other way around.

- Ask for your level — first-year-PhD legibility.
- Ask for your form — examples, figures, equations, pseudocode.
- Ask for your notation and chain of logic — not generic modelese.
- Iterate until you can teach the result in your own words.

The agent works for the researcher, not the other way around.

COMMANDMENT II

AI is not your peer.
It is your cyber-arm.

Do not reserve AI for tasks you dislike. Attach it to the work you care about most.

A peer is someone you hand tasks to. A cyber-arm is part of your thinking loop — while you reason, write, formulate, rehearse, decide, and worry.

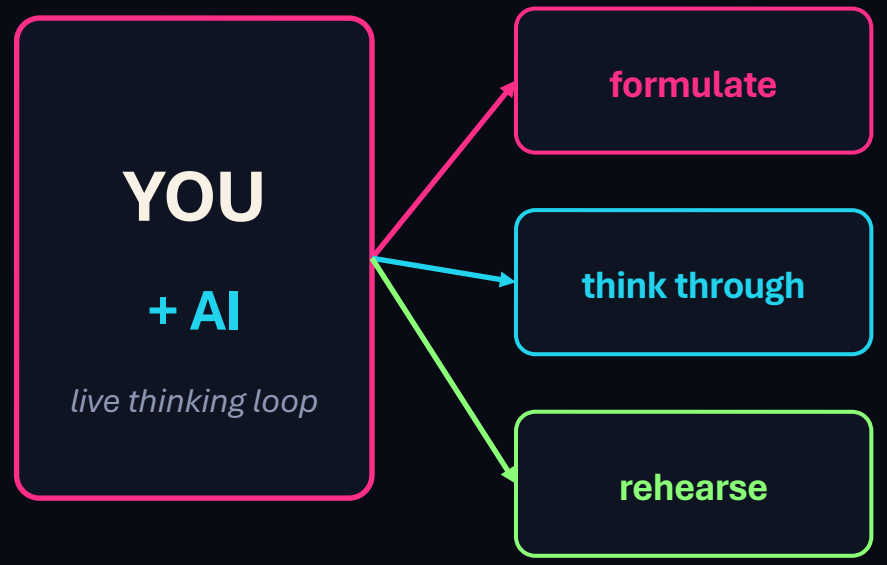

CORE REFLEX

If an idea is active in your mind, put it into the AI loop before it cools down.

Do not ask only for the chores.
Ask hardest where you are most invested.

Whatever concerns you, have AI think about it with you.

OPERATING POSTURE

Turn outputs into owned knowledge.

Do not accept eloquence as understanding. Push the agent until the output enters your own mental model.

01	**Simplify**	Explain it at first-year PhD level; define all terms.
02	**Illustrate**	Make a figure, an analogy, an example, and a counterexample.
03	**Interrogate**	Show evidence, assumptions, weak points, and alternatives.
04	**Teach back**	Rewrite in your own words until you can present it.

USEFUL PROMPTS

“Make it simpler.”

“Give me a concrete example.”

“Draw the mechanism.”

“What evidence supports this?”

“Argue against your own answer.”

Understanding is the gate between machine output and human knowledge.

COMMANDMENT III

Our input–output bandwidth is the bottleneck.

Agent power is abundant. Human specification and comprehension are scarce.

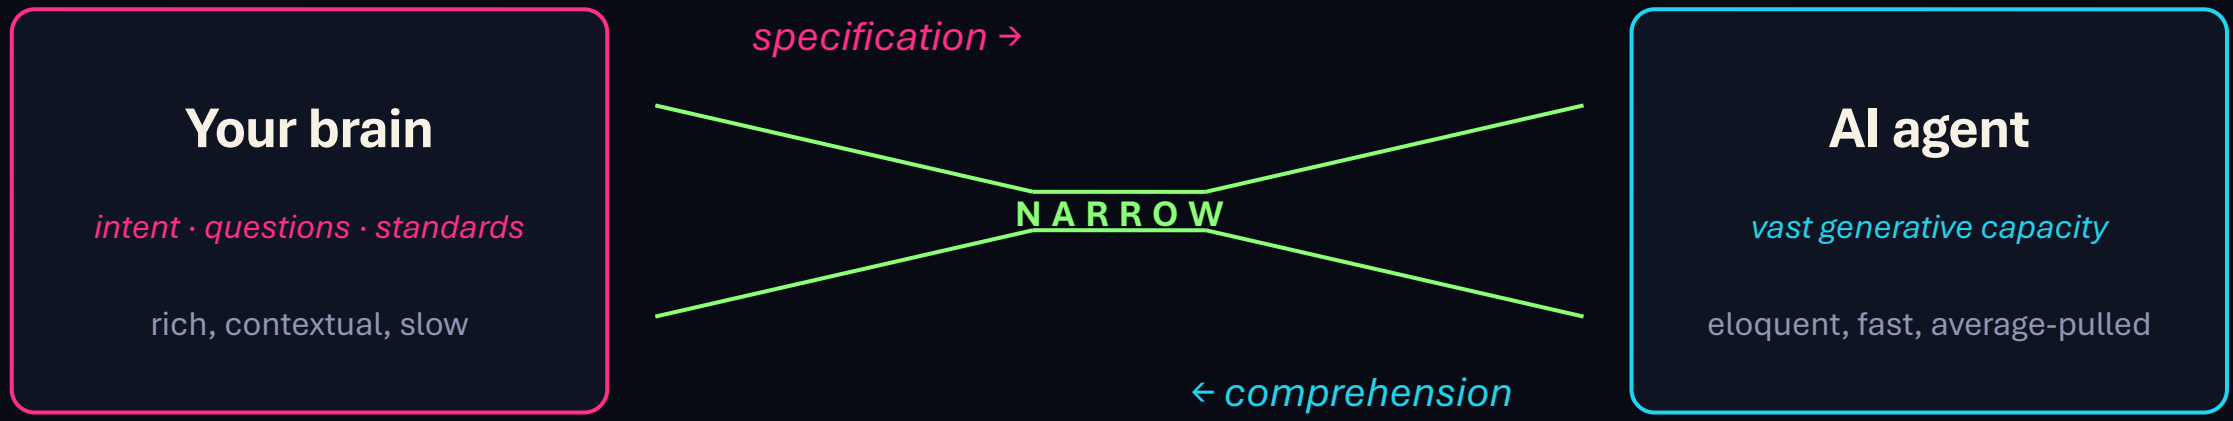

Not the model's power — our bandwidth.

Without specification, agents regress to average.
Without comprehension, you lose the steering wheel.

The real research skill is moving more of your intent into text, and more of the agent's output into your mind.

If the channel is narrow, widen the channel — do not blame the ocean.

PRACTICAL PROTOCOL

Braindump before delegation.

Write a separate spec file. Pour your whole research state into it — not polished prose, raw intent.

WHAT GOES IN THE SPEC

Goal — what would count as a useful result
Context — the story or logic already in your head
Constraints — audience, style, methods, forbidden routes
Known unknowns — where you are uncertain or confused
Desired output — format, depth, examples, figures, tests
Validation — what evidence would make you trust it

Starter instruction: *"Read spec.md. Ask only blocking questions. Then produce version 0, with assumptions marked."*

A good spec pulls the agent toward your corner of the problem space.

PCAPS — chain of logic

P **Problem**
what are you tackling?

C **Challenge**
why is it not yet solved?

A **Approach**
what is your take, and why?

P **Proof**
what evidence shows it works?

S **Significance**
the contribution — and who cares?

COMMANDMENT IV

Be above AI on everything you think about.

Generic AI output is the floor, not the ceiling.

Before claiming competence on any research activity, you should at least know what a generic AI agent can say, plan, and explain from your inputs — and be able to communicate it yourself.

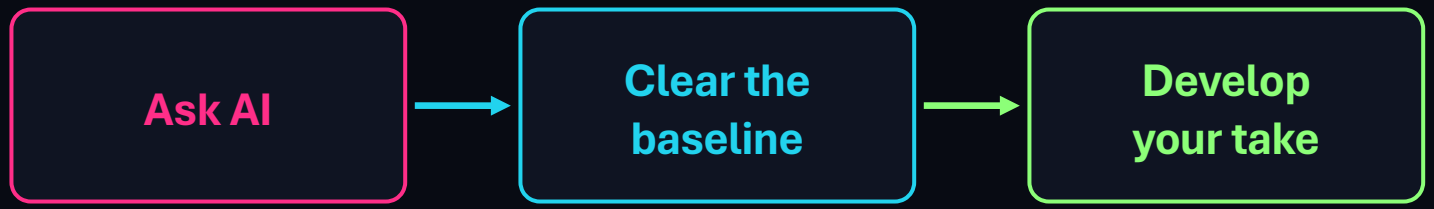

MINIMUM STANDARD

- Writing a story? Know and beat the AI version.
- Entering a field? Understand the concepts and works AI surfaces.
- Planning an experiment? Plan at least as well as generic AI from what you know.
- Only then are you ready to go beyond the generic answer.

Below generic AI is below the minimum standard.

Use AI to establish the baseline; use your research to exceed it.

THE WORKFLOW

Four commandments. One loop.

For every serious research task with AI, cycle until you can defend the result without hiding behind the model.

1	**Own the output**	Explain, justify, simplify, and teach back.
2	**Attach to live thought**	Bounce active ideas, presentations, formulations, and concerns off the agent.
3	**Expand bandwidth**	Braindump goals, logic, constraints, audience, and tests.
4	**Clear the AI baseline**	Know the generic answer before claiming your own expertise.
5	**Go beyond it**	Make your original move: sharper evidence, better framing, stronger judgment.

FINAL STANDARD

"I can explain it. I have used AI on it. I can beat the generic answer. I can defend my own take."

Research-with-AI is delegation under ownership, attached to thought, above the baseline.

PART II

The tricks.

The workflow: seven habits for research with AI agents, distilled from practice.

1	**Keep a hierarchy of files.**	*read at any zoom level*
2	**Iterate the spine, not the details.**	*feedback where it lasts*
3	**Never be afraid to start afresh.**	*restart from idea.md*
4	**Ask for snippets, not finished slides.**	*you compose, agents draw*
5	**Fix a palette and a plot style once.**	*a recognizable look*
6	**Keep a local library of full texts.**	*agents flag, you fetch*
7	**Keep the context window fresh.**	*compact at checkpoints*

Tricks change with the tools. The stance does not.

TRICK 1 · FILE HIERARCHY

Keep a hierarchy of files.

One project, readable at three zoom levels.

Read at the level your time allows: skim the flyer, work the spine, or dive into the full text. Steer from whichever level you are on.

1-page flyer · *the story at a glance*

10-page spine · *the core chain of logic*

100-page manuscript · *every detail, fully expanded*

Perspective files.

Same content, viewed from different angles.

- manuscript-design.md: the design you intend for the piece.
- results-background.md: methods and concepts explained until you fully understand.
- Each view lets you audit a different aspect of the work.
- Agents keep every view in sync.

Read at the zoom you have time for. Agents keep the levels aligned.

TRICK 2 · FEEDBACK ALTITUDE

Iterate the spine, not the details.

Human feedback belongs at the level of story and logic.

Details churn: agents rewrite or revert them across rounds. The spine lasts: the narrative, the chain of logic, the direction. Feedback there compounds.

YOU ITERATE

Pivot the narrative.

Rework the chain of logic.

Point at new directions.

AGENTS ITERATE

Rewrite wording and sections.

Rerun derivations, code, figures.

Keep every file consistent.

Exception: *polish the details only when preparing to publish, or to show your work to others.*

Speak one level above the agents. Let them churn below you.

TRICK 3 · FRESH STARTS

Never be afraid to start afresh.

Restart from understanding, not from the pile.

Projects accumulate manuscripts, code, results, side tracks. Much of it is redundant or half-analyzed. By now your head holds a clearer story than the folder does.

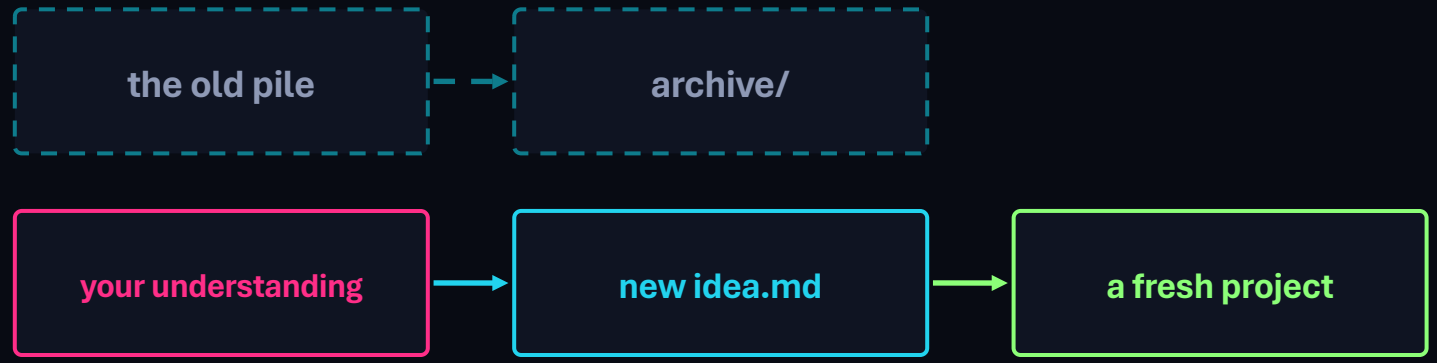

The move.

Migrate your understanding. Nothing else.

- Everything so far goes into archive/.
- Write a new idea.md: only what is distilled in your brain.
- Start the project again from that file alone.
- Agents may dig into the archive, but only when you say the word.

A fresh start from a clear story beats another round of patching.

TRICK 4 · FIGURES & SLIDES

Ask for snippets, not finished slides.

Agents draw the parts. You compose the whole.

Iconic panels and image snippets are easy to specify and easy to judge. Position, size, and animation are quick by hand, and slow to dictate in words.

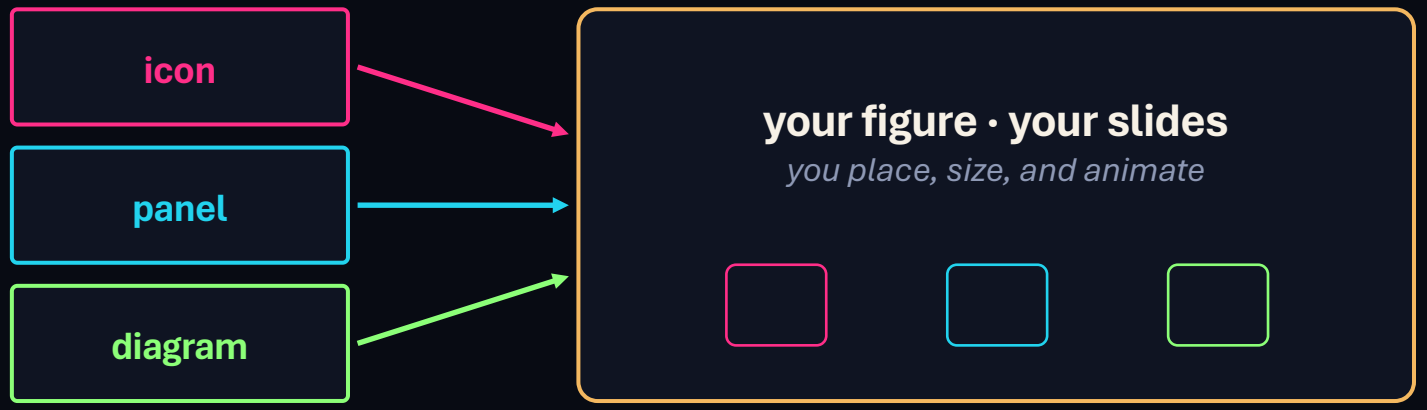

Why it works.

Small asks, sharp judgments.

- A snippet is one visual idea: easy to prompt, easy to iterate.
- Layout is spatial: your hands beat your words.
- The same snippet serves paper, talk, and poster.

Keep visual control. Delegate visual components.

TRICK 5 · PLOT STYLE

Fix a palette and a plot style once.

Every figure lands in the same visual language.

Set the palette and plot defaults in one style file at the repo root, applied on import. Point every agent session at it.

one palette, fixed order, everywhere

Same colors in the paper, the talk, and the poster.
Match the figure fonts to the document fonts, too.

Payoff.

Consistency compounds.

- No per-figure styling decisions.
- Figures from different sessions still match.
- One import line applies the whole style.
- Your work becomes recognizable at a glance.

A consistent look across every figure, talk, and paper.

TRICK 6 · LITERATURE

Keep a local library of full texts.

Agents know abstracts. Research runs on full texts.

Model memory holds mostly abstracts. Web search misses paywalled papers with no open preprint. A folder of full-text PDFs closes the gap.

STANDING INSTRUCTION

“While you work, keep a running list of the papers whose full text would most improve the result. Tell me the top ones. I will go fetch them.”

You have library access. The agents do not. Be their fetch arm.

TRICK 7 · CONTEXT

Keep the context window fresh.

Stay below 300k tokens. Compact at clean checkpoints.

Attention and quality degrade as the context fills. Do not run to the hard limit and get cut mid-thought.

01	**Watch**	Check context usage as each phase of work completes.
02	**Pause**	Above 300k, ask the agent to pause at a clean checkpoint: end of a phase, a compilable state.
03	**Compact**	Compress the context, carry the essentials, keep going.

Pause at checkpoints, not mid-thought.

THE TOOLKIT

Seven tricks. One toolkit.

The commandments set the stance. These habits make it run, day to day.

1	**Hierarchy of files**	Three zoom levels, plus perspective files, kept aligned.
2	**Spine over details**	Feedback on story and logic. Agents churn the rest.
3	**Fresh starts**	Archive the pile. Restart from a distilled idea.md.
4	**Snippets, not slides**	Agents draw the components. You compose and animate.
5	**One style**	A fixed palette and plot defaults, applied everywhere.
6	**Local library**	Agents flag the papers. You fetch the full texts.
7	**Fresh context**	Below 300k. Compact at clean checkpoints.

WORKING STANCE

"Agents supply the volume. I keep the structure, the story, and the standards."

Commandments set the stance. Tricks make it run.