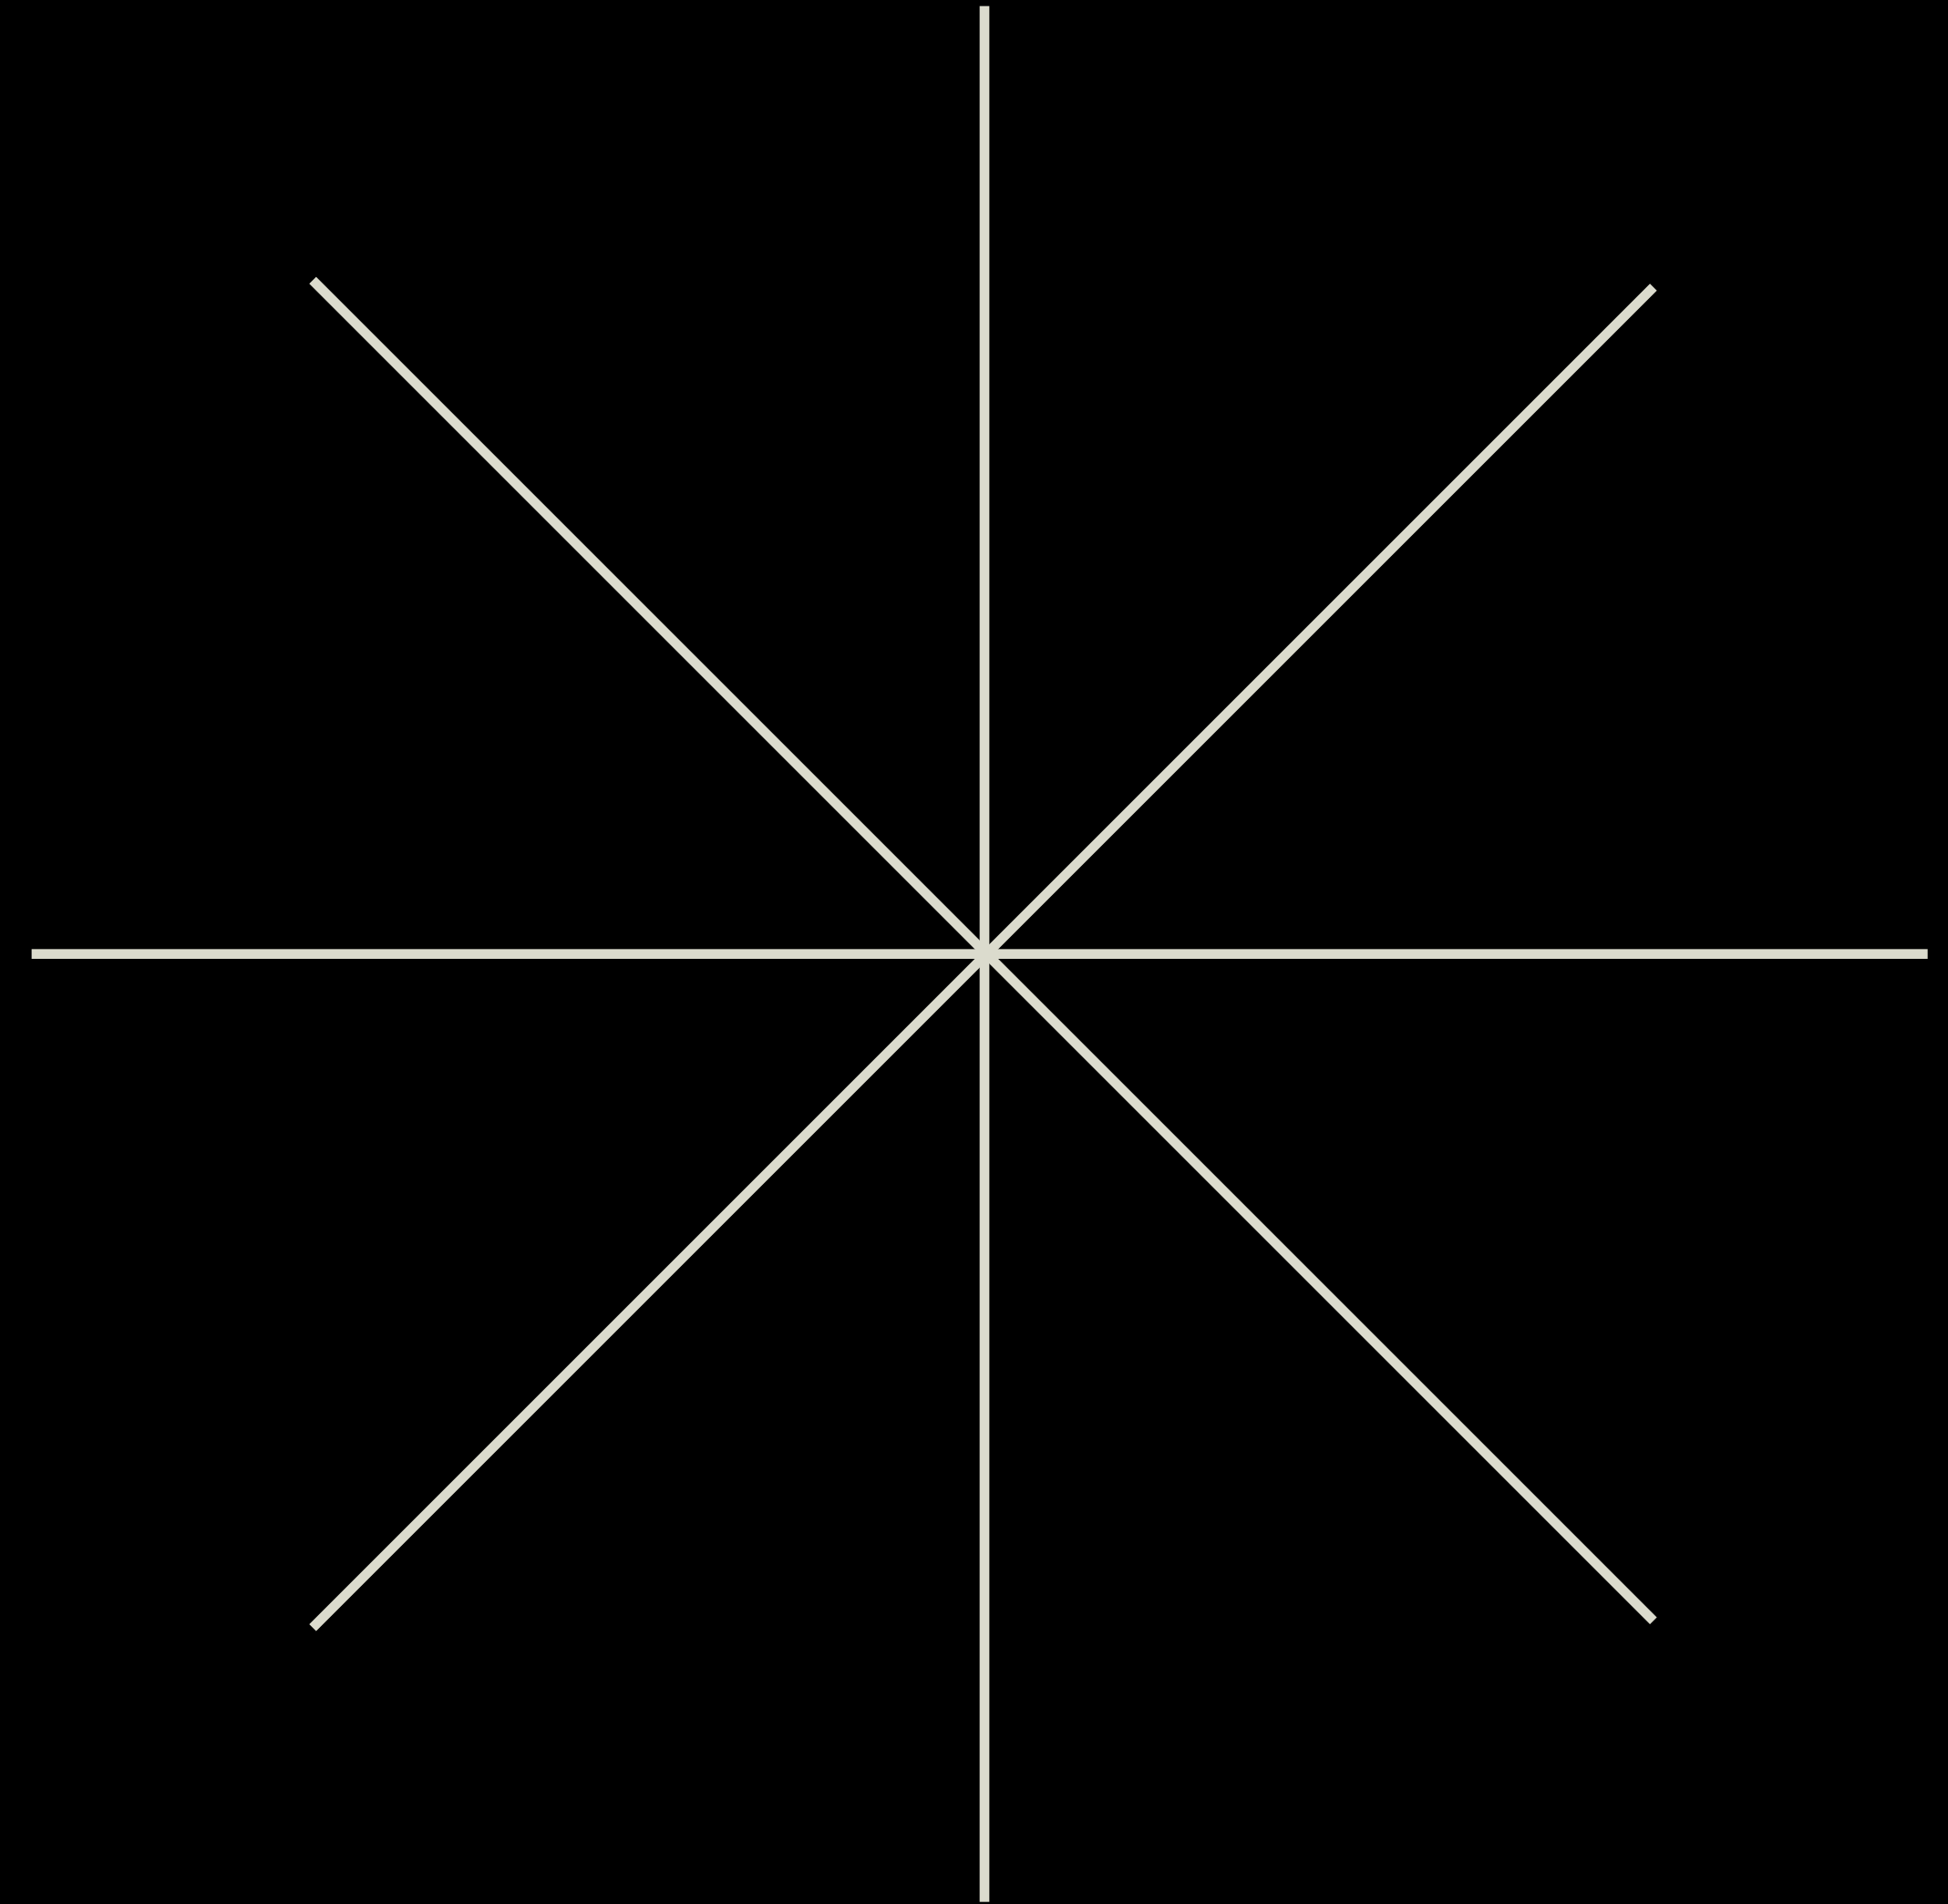
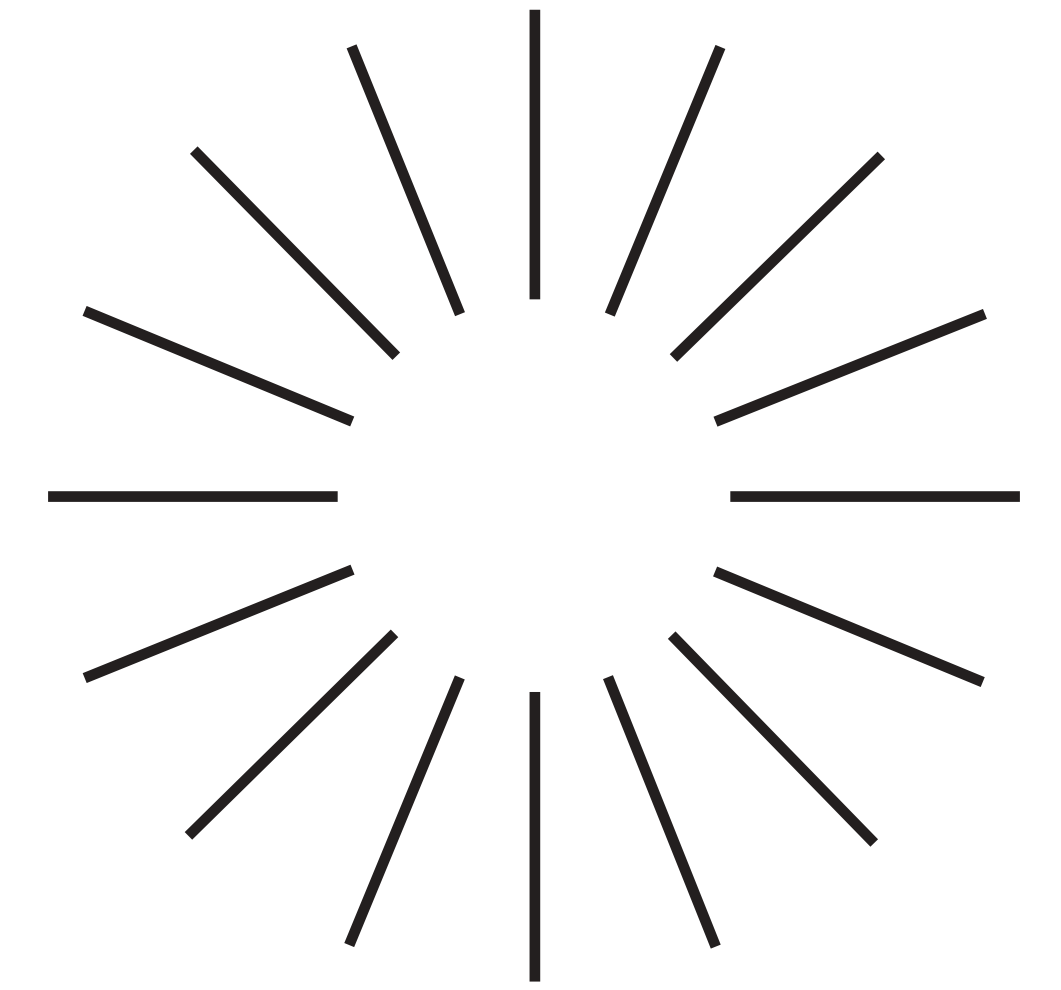


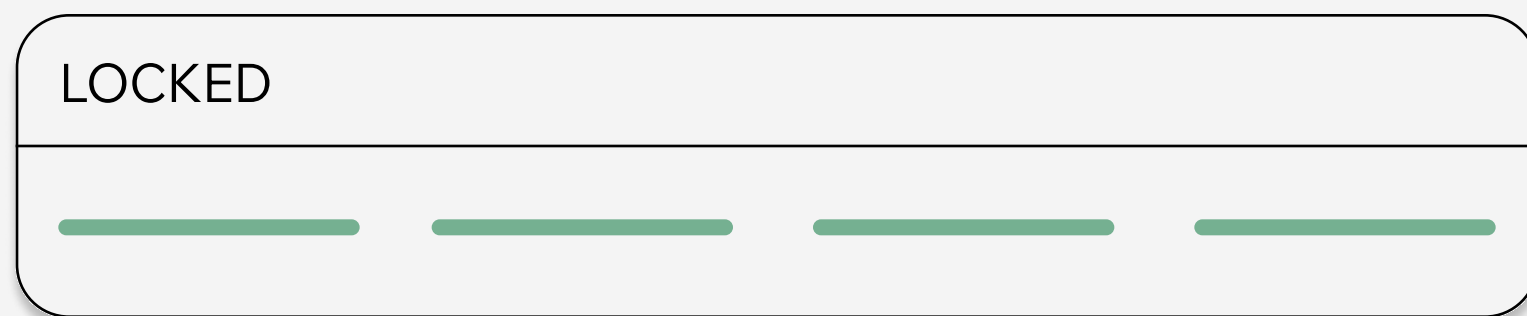
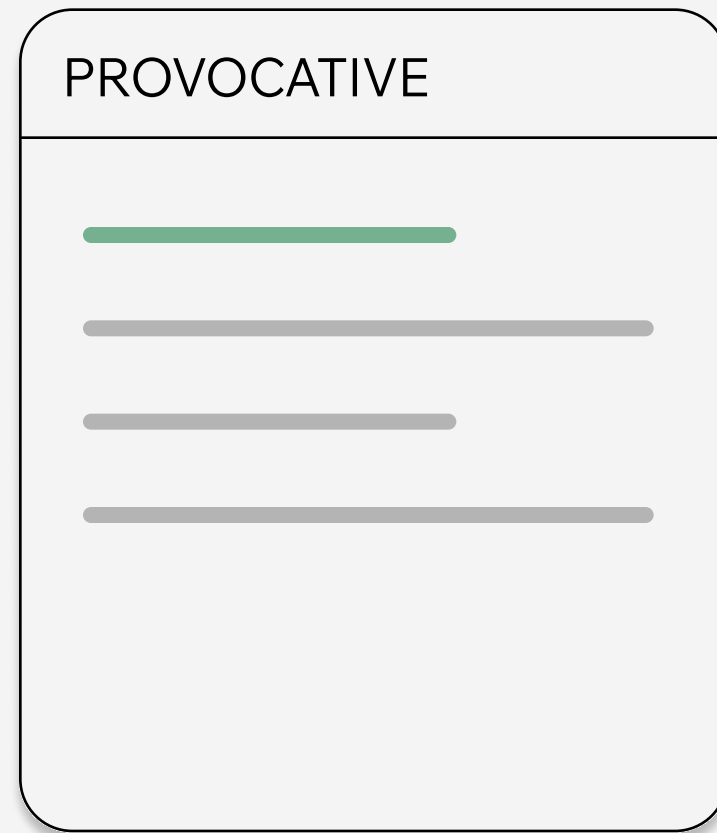
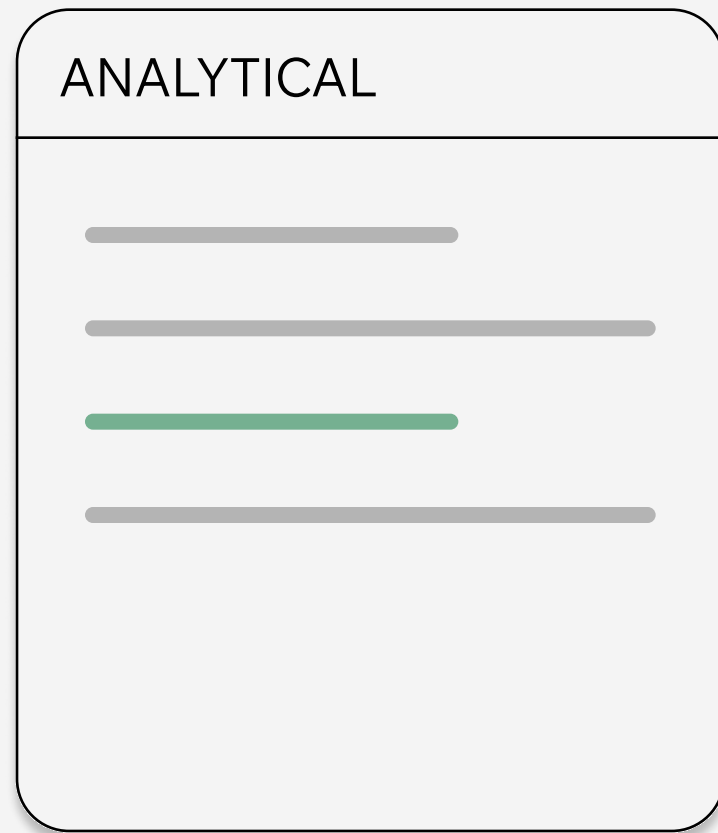
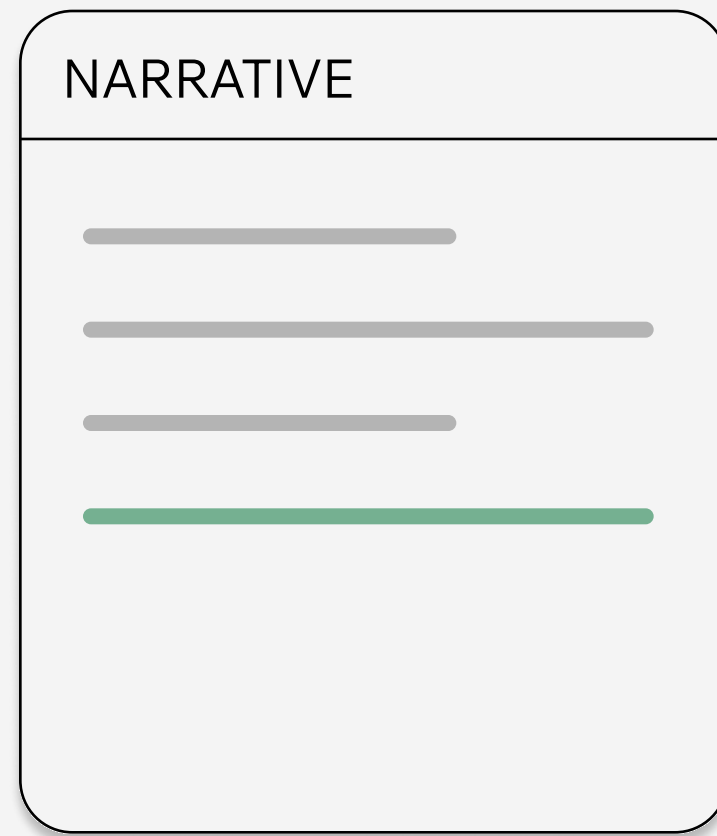
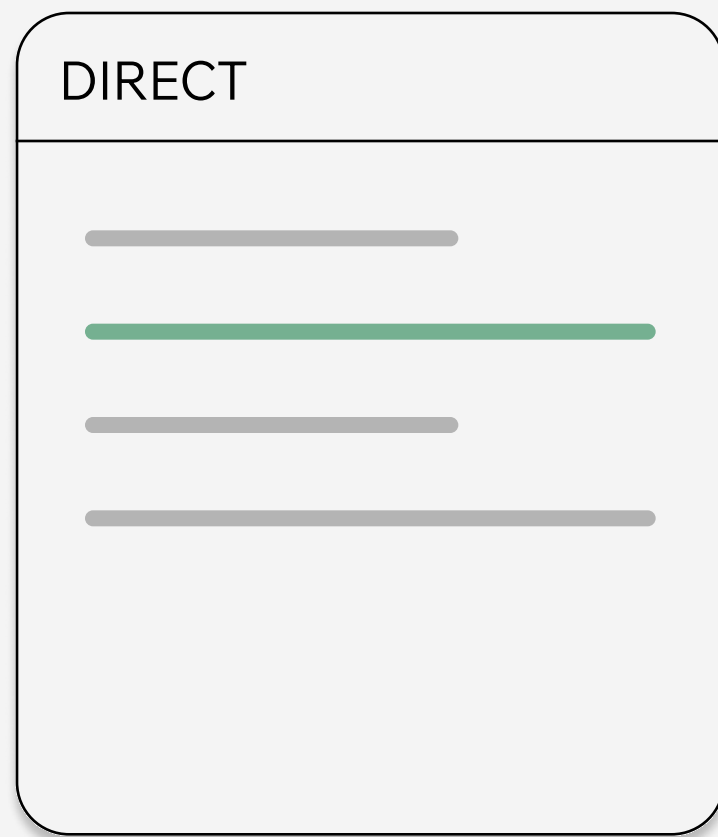
FINAL PROJECT



AGENDA

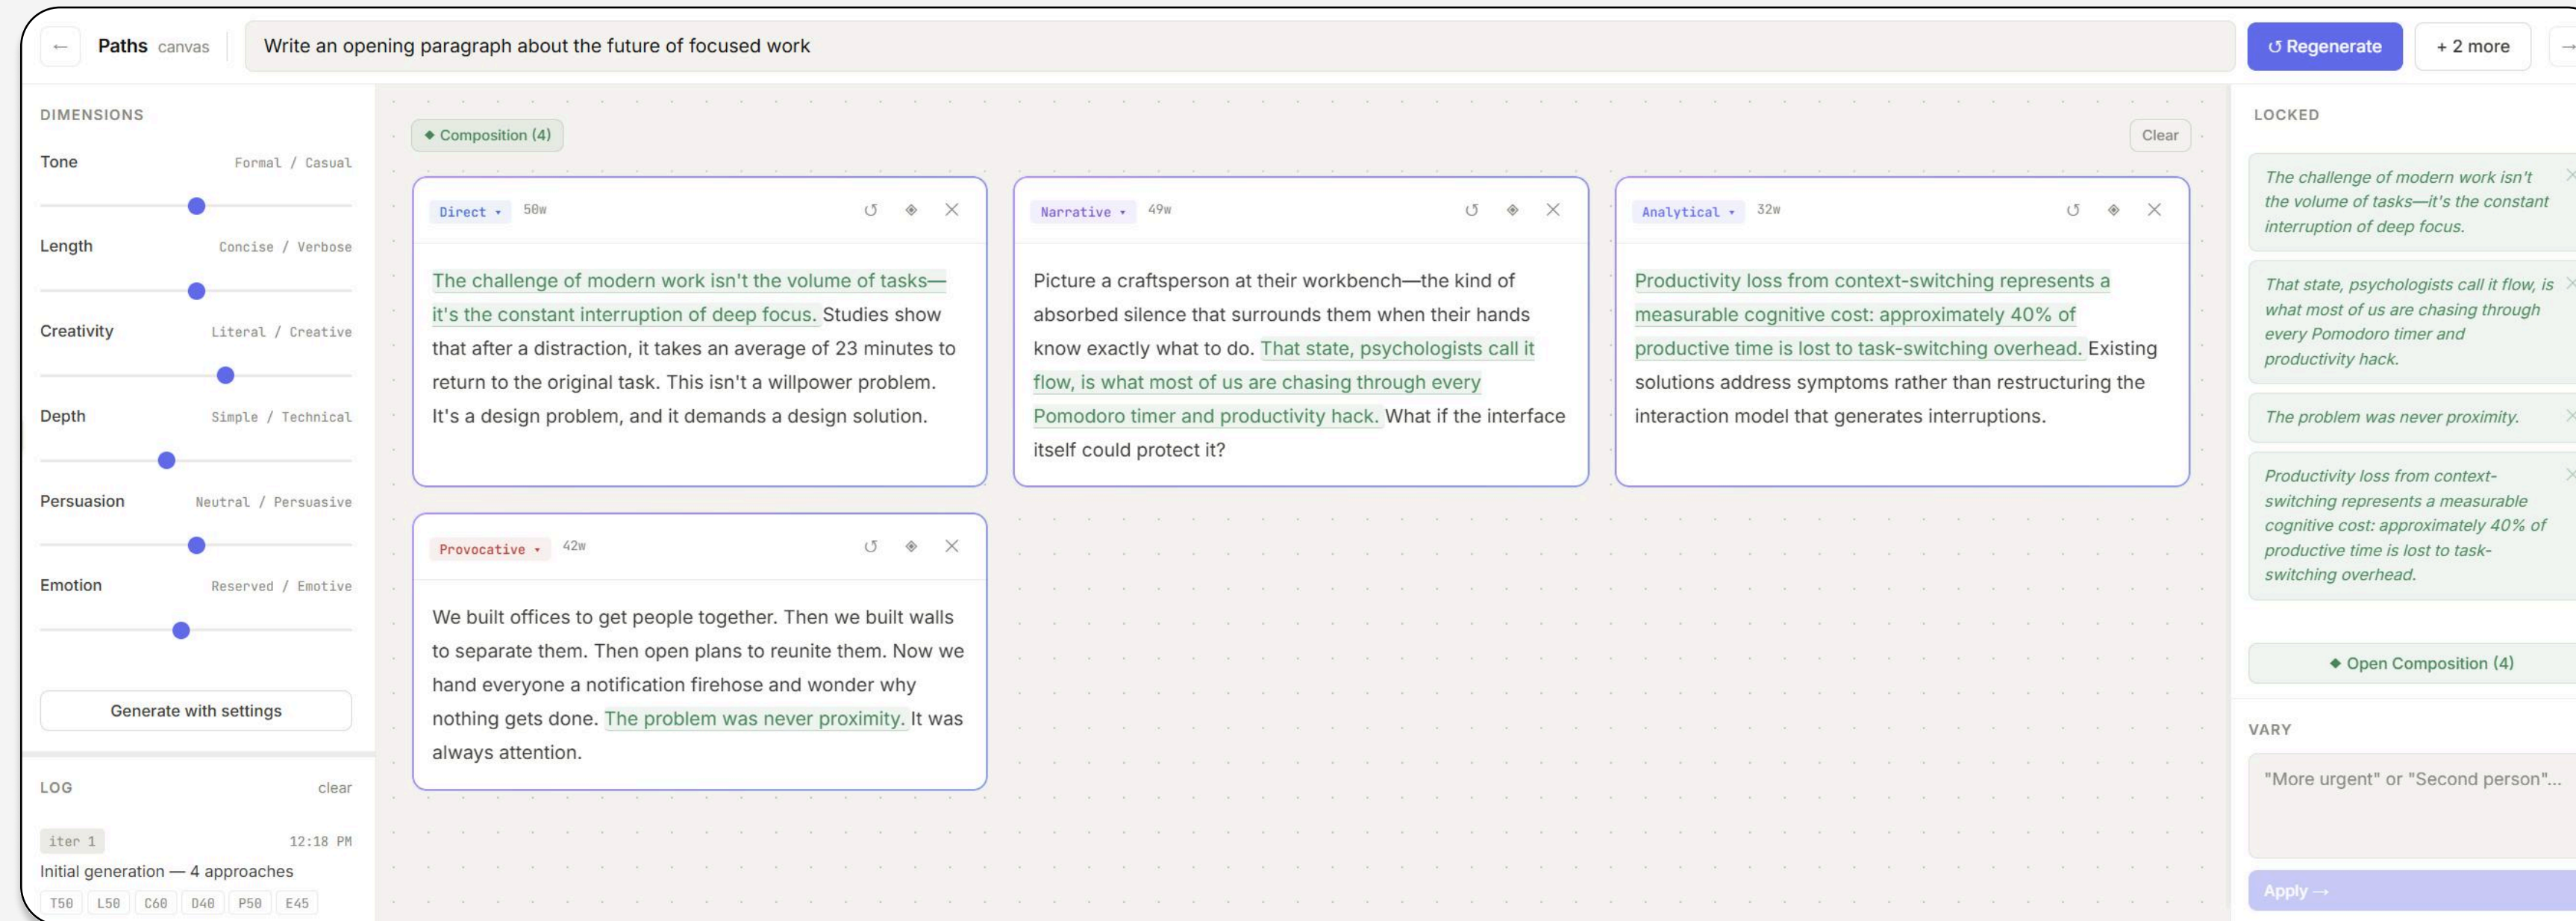


- 1—What & Why
- 2—How It Works
- 3—Key Decisions
- 4—Observed Behaviors
- 5—Reflection



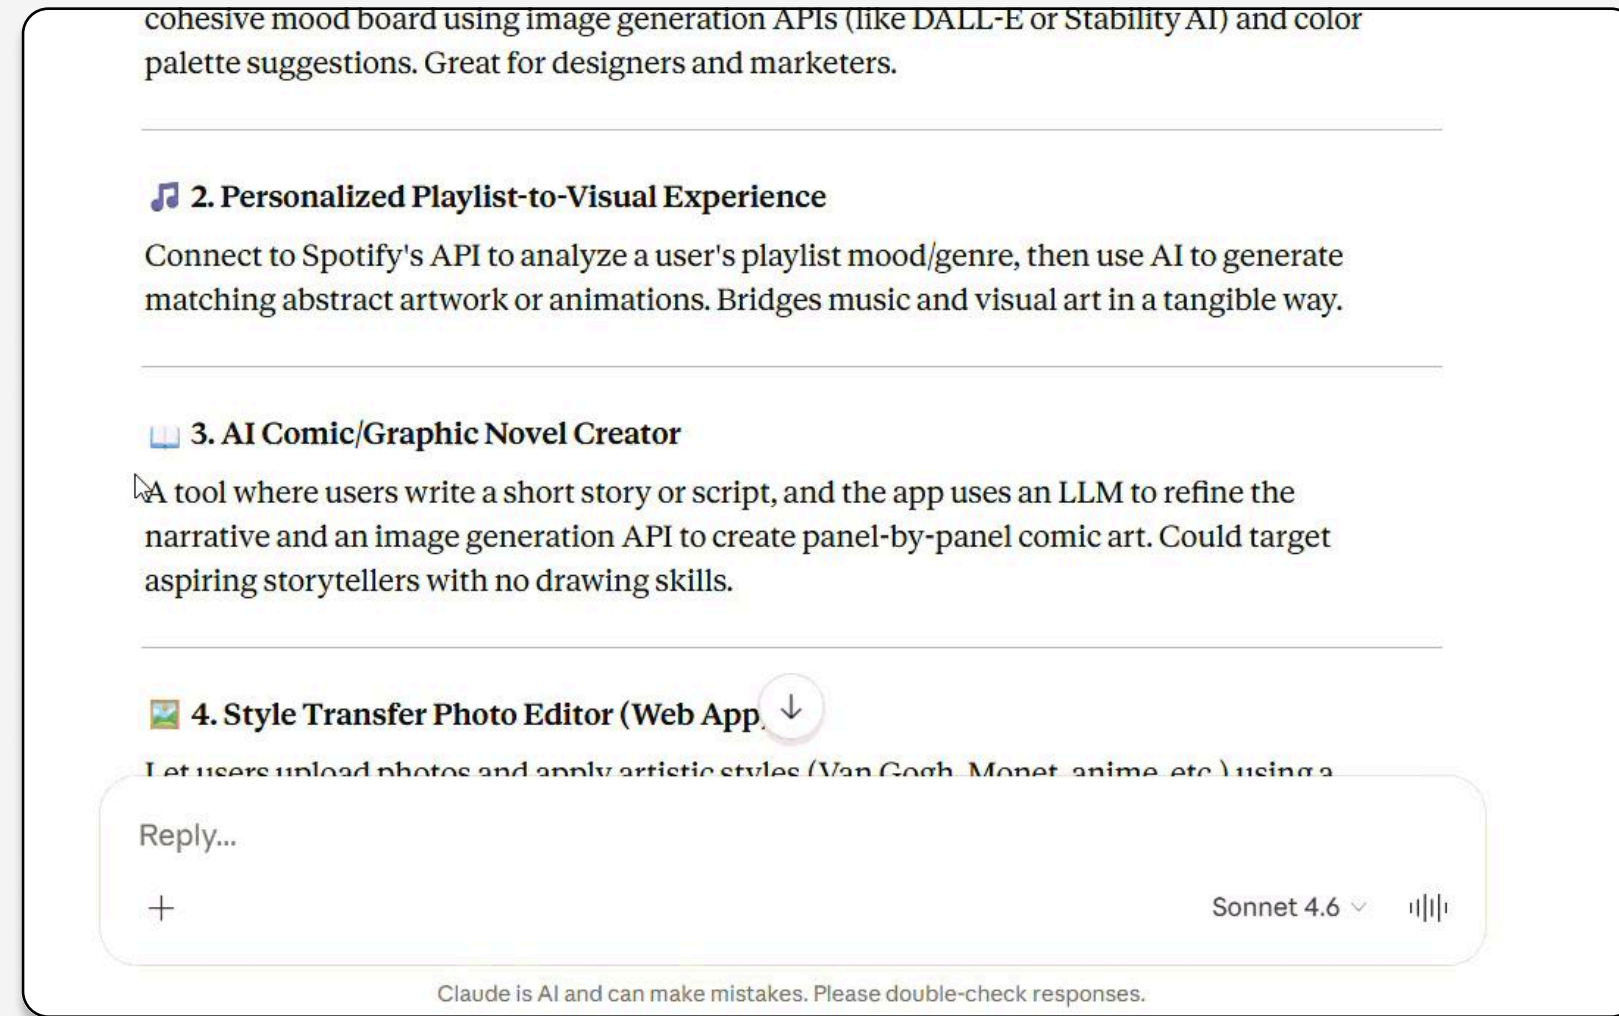
PATHS

[LINK TO PROTOTYPE](#)



A canvas based generative writing tool that lets users explore, compare, and lock AI-generated text variations.

PATHS



The problem with AI writing tools

LINEAR CHAT

The Problem

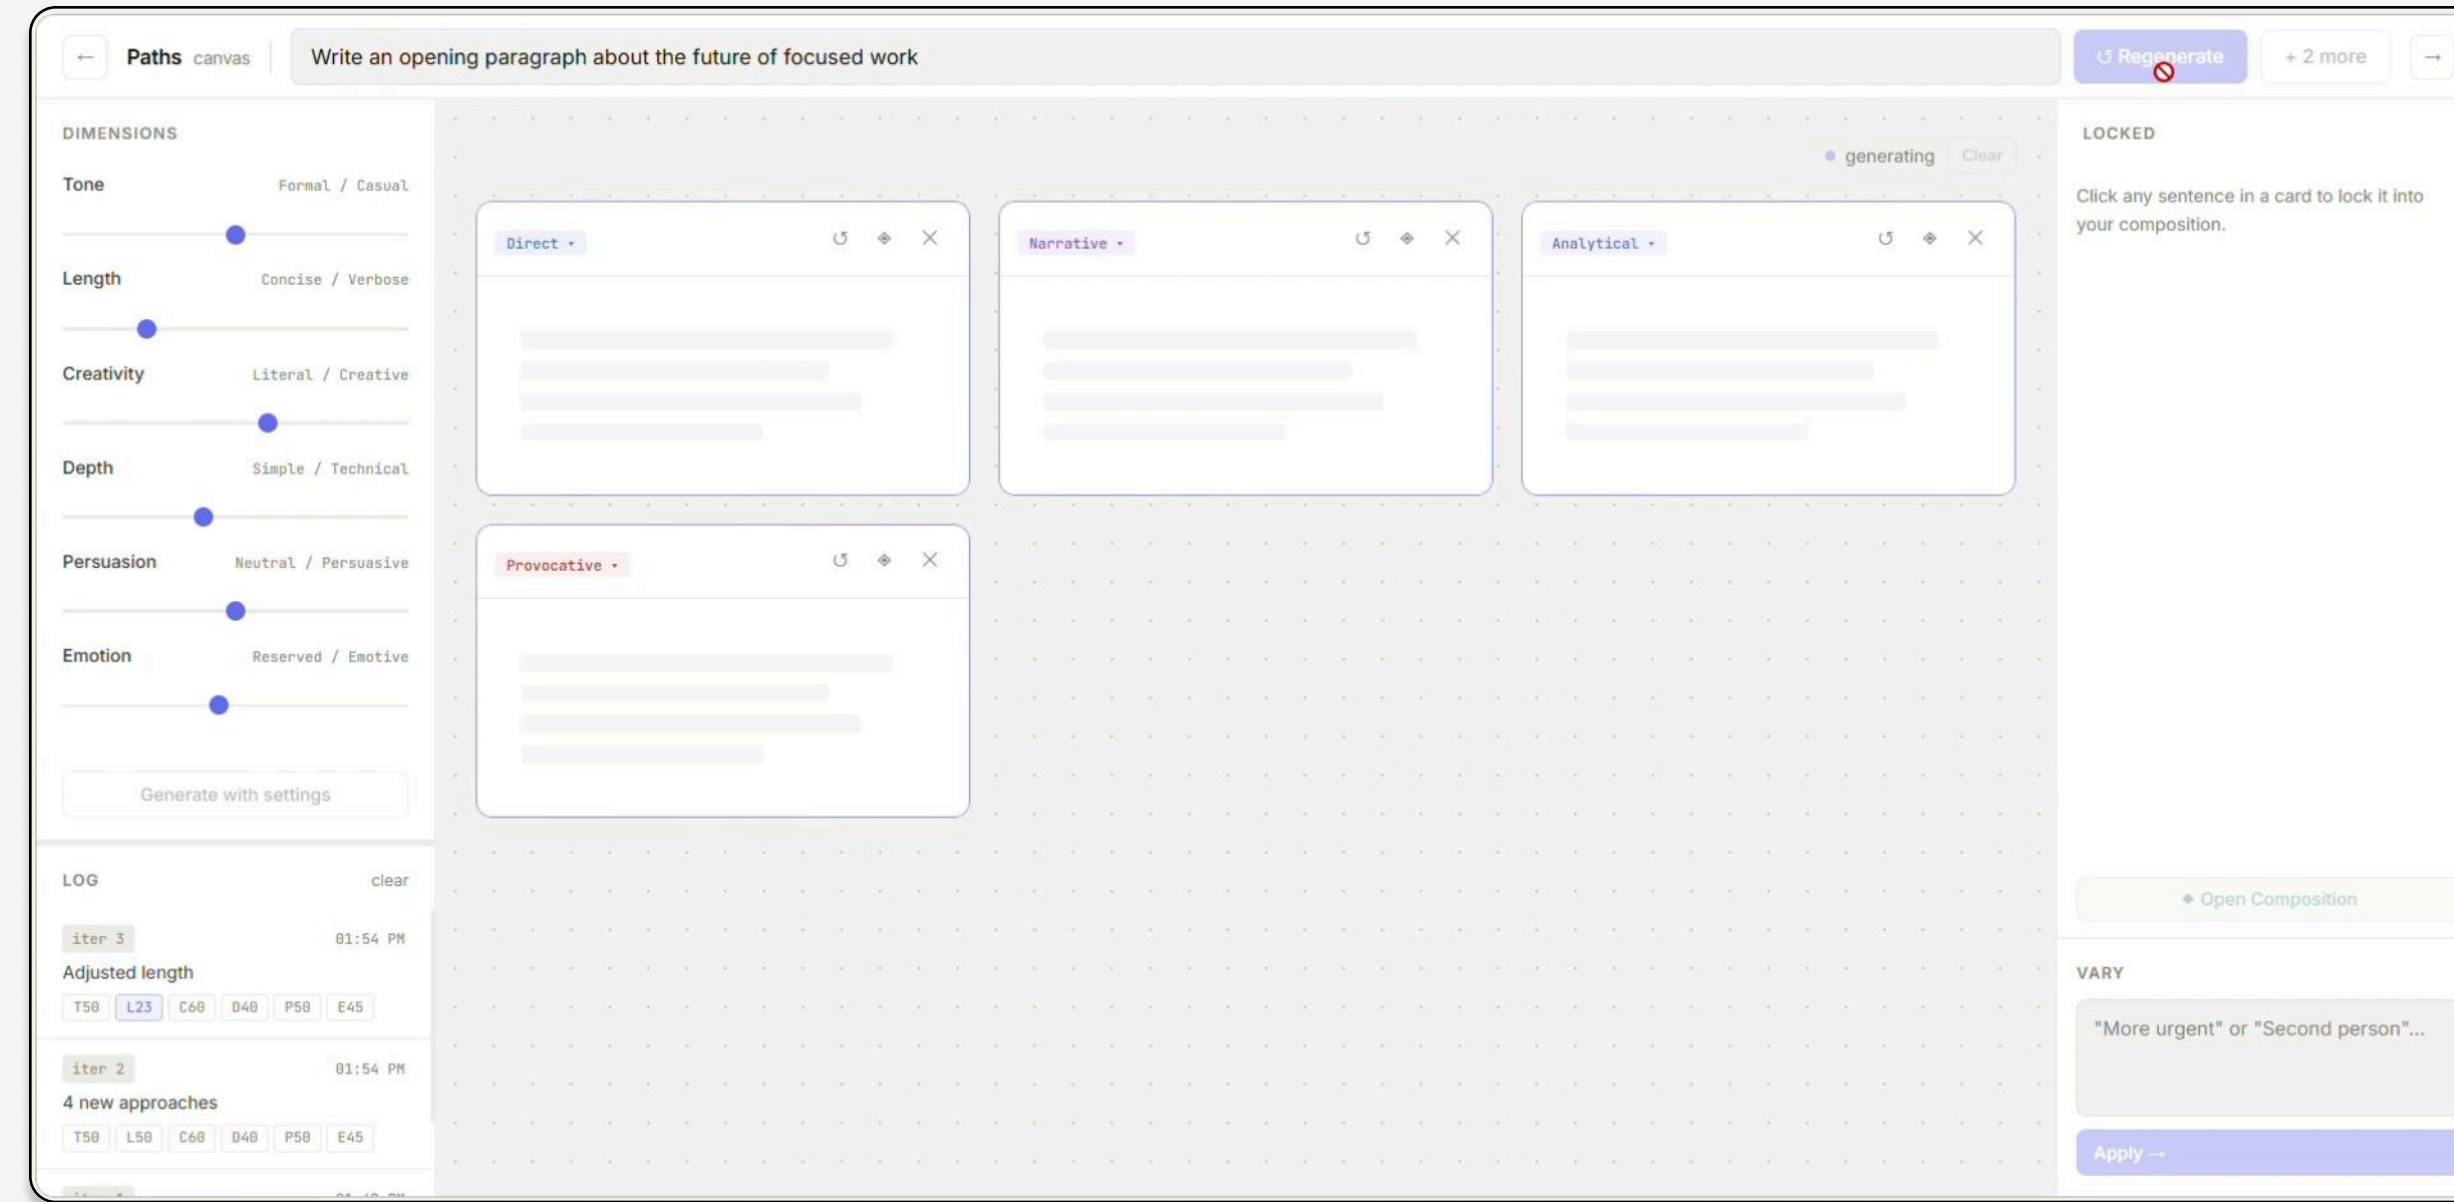
Most AI writing tools generate one response at a time in a linear chat. Users can't compare alternatives, explore interpretations, or preserve the parts they like while changing the rest.

VARIATION

The Solution

Paths generates multiple style variations side-by-side. Users can lock sentences they want to keep, regenerate everything else, and compose a final piece from their favorite fragments.

PATHS



WHO IS IT FOR?

Writers, communicators, and anyone who uses AI to draft content.

PURPOSE

Iterative creative exploration:

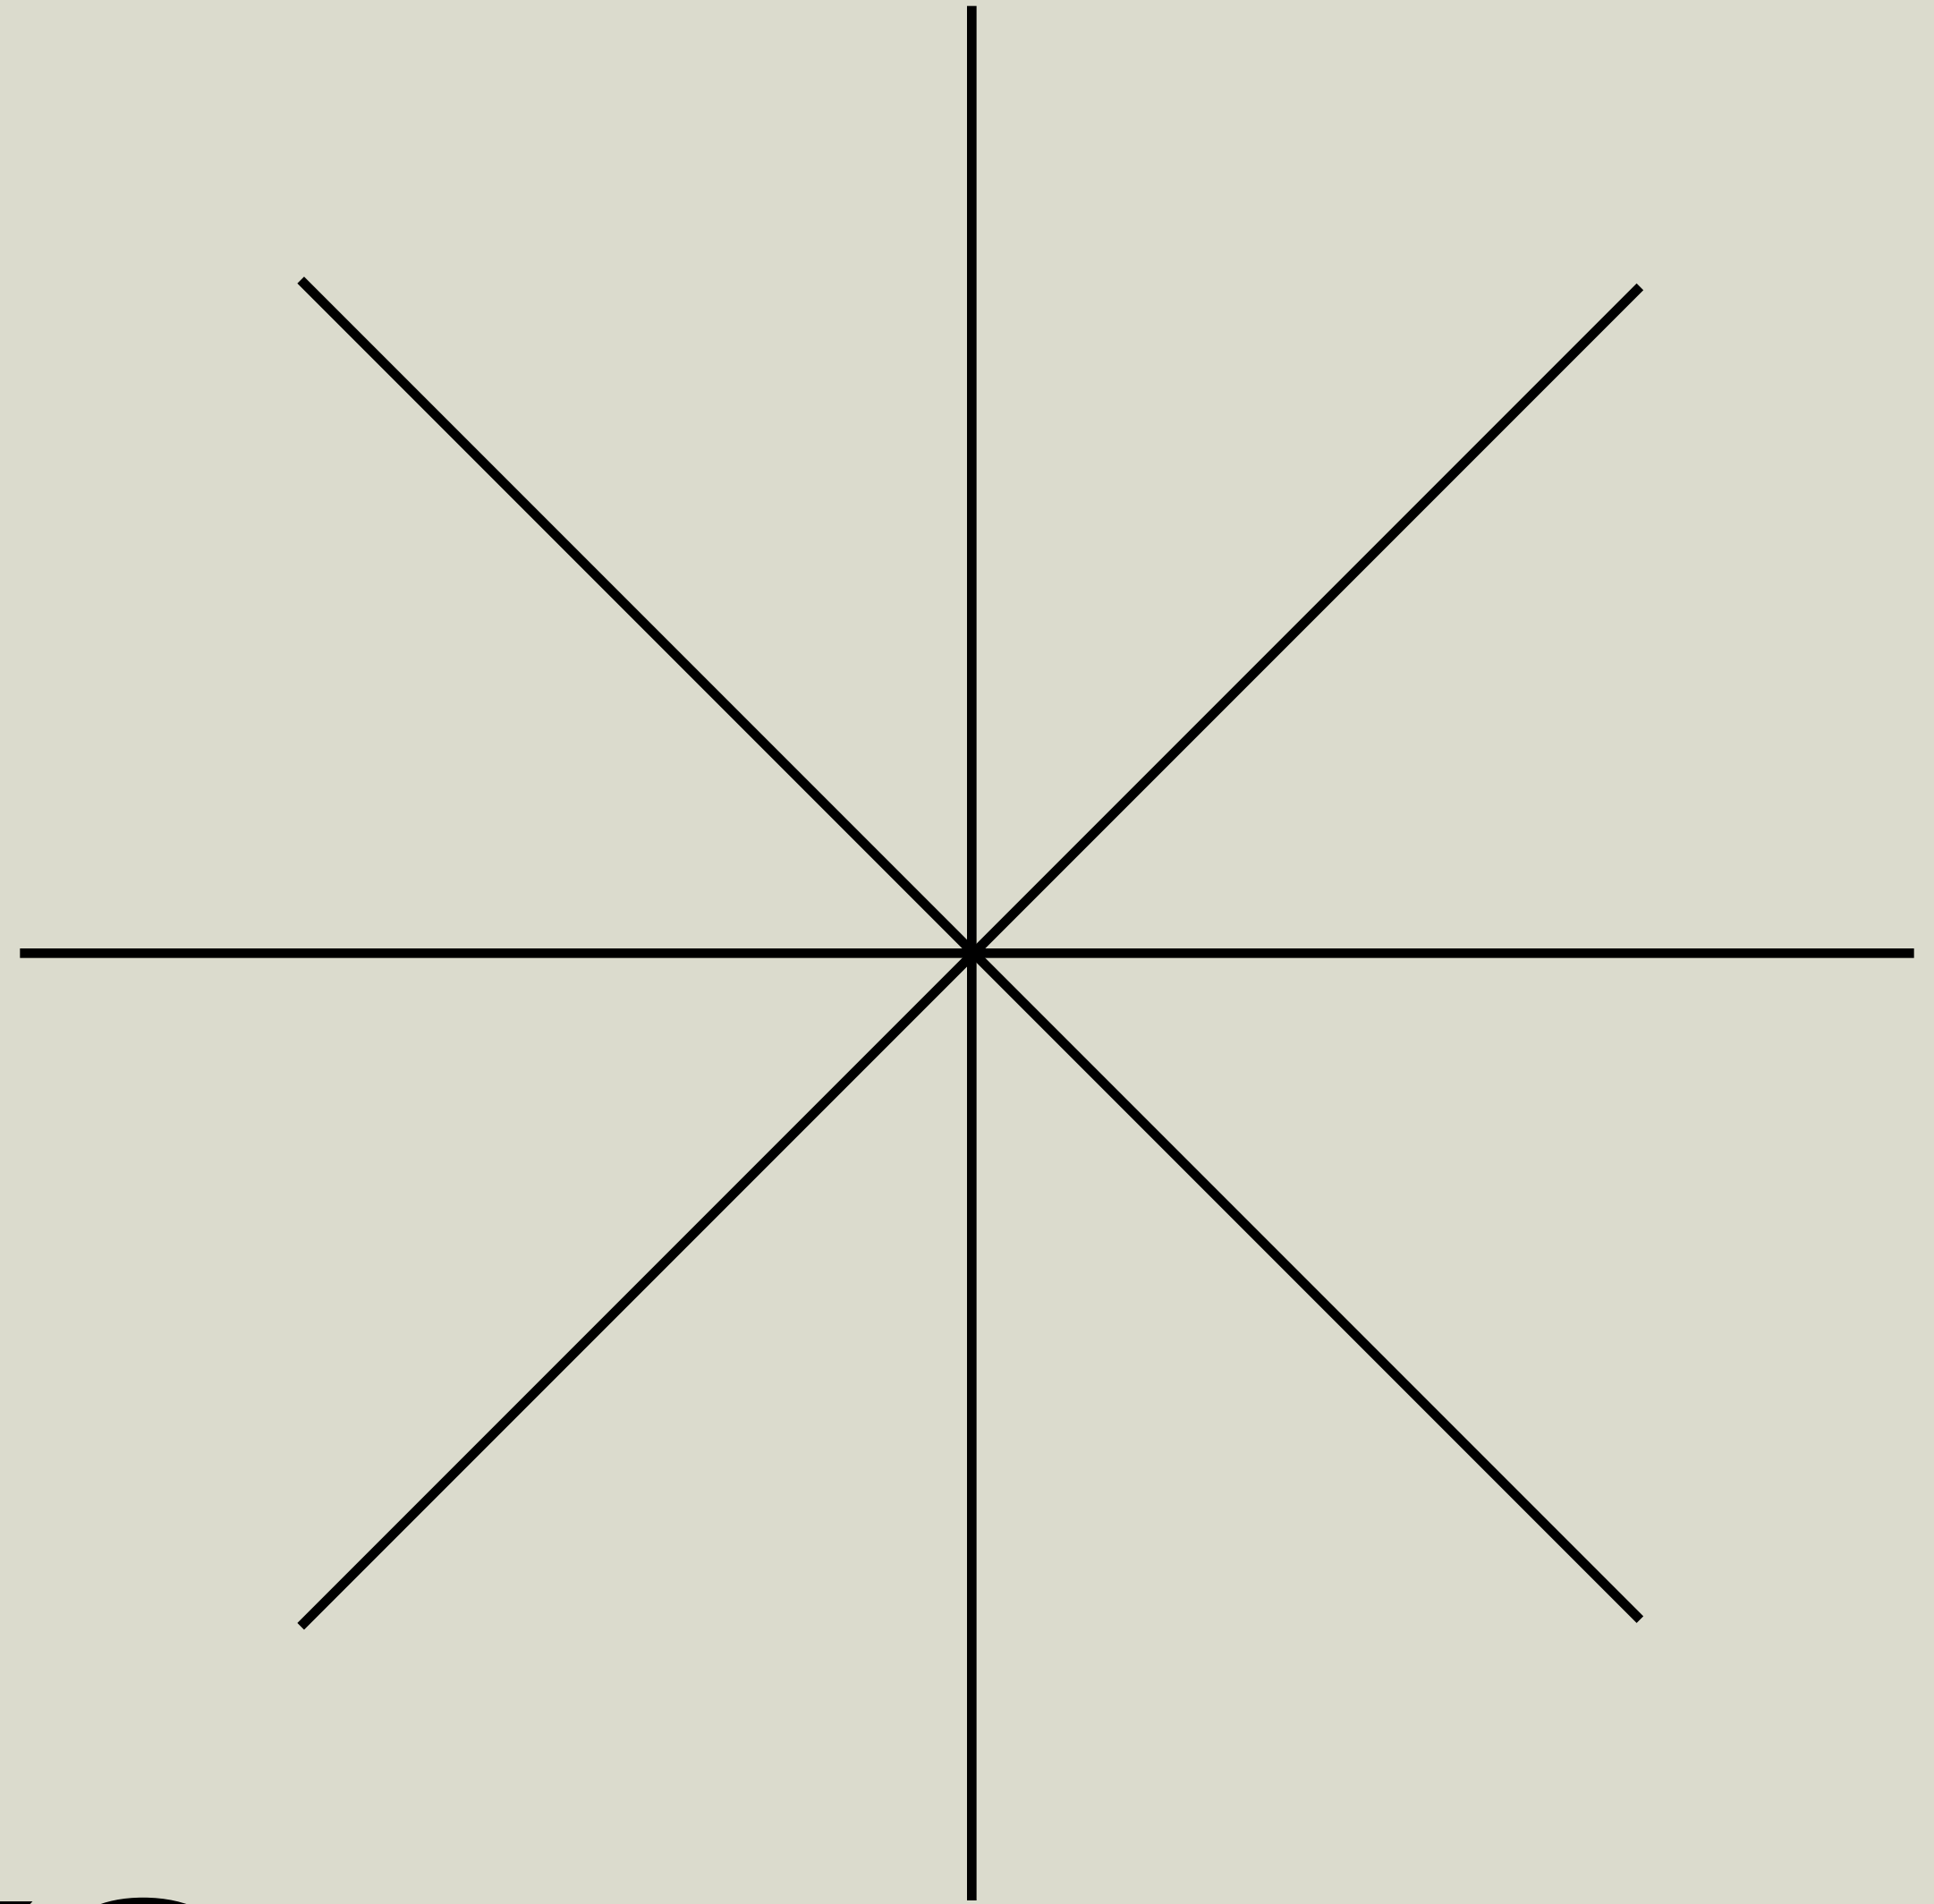
utility + expression

CONTEXT

Early-stage writing, brainstorming, or refining tone and voice.

02

HOW IT WORKS



From prompt to composition

1

Enter a prompt

Users type a writing task into the top bar.

Dimension sliders set the generation parameters.

2

Generate variations

Paths generates by default 4 cards simultaneously using the same prompt but different style instructions.

3

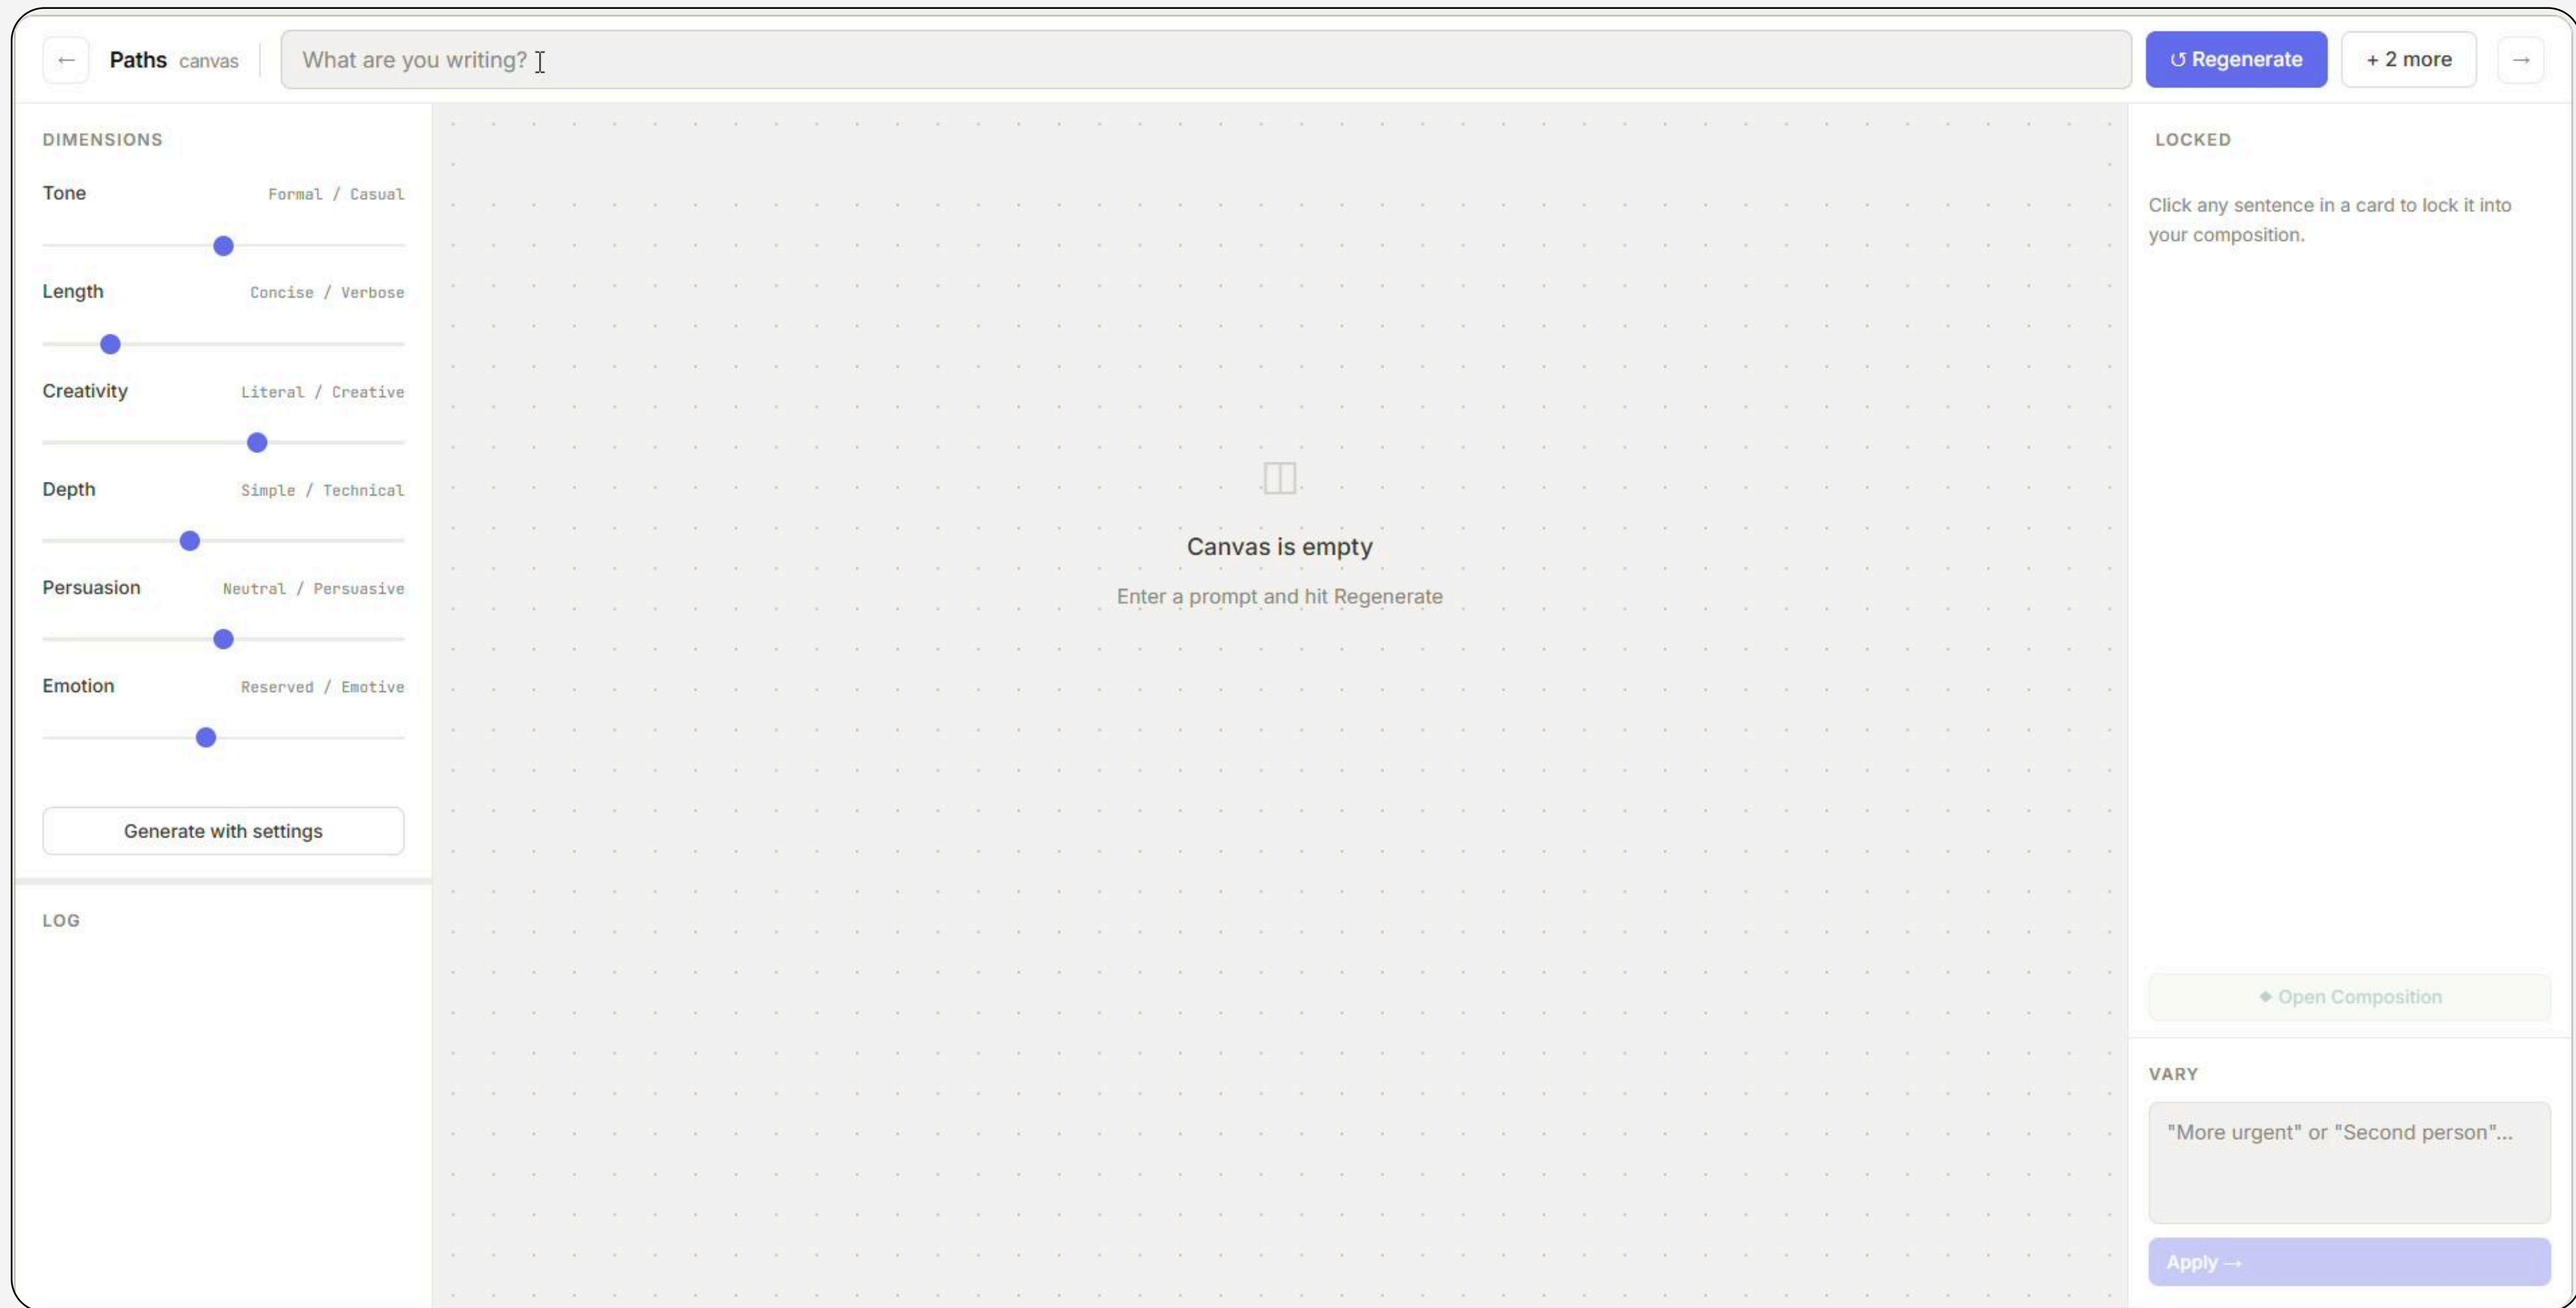
Lock sentences

Users hover over any sentence and click to lock it. Locked sentences appear in the right panel and are passed back to the model as 'Include verbatim: [constraints]'

4

Regenerate & compose

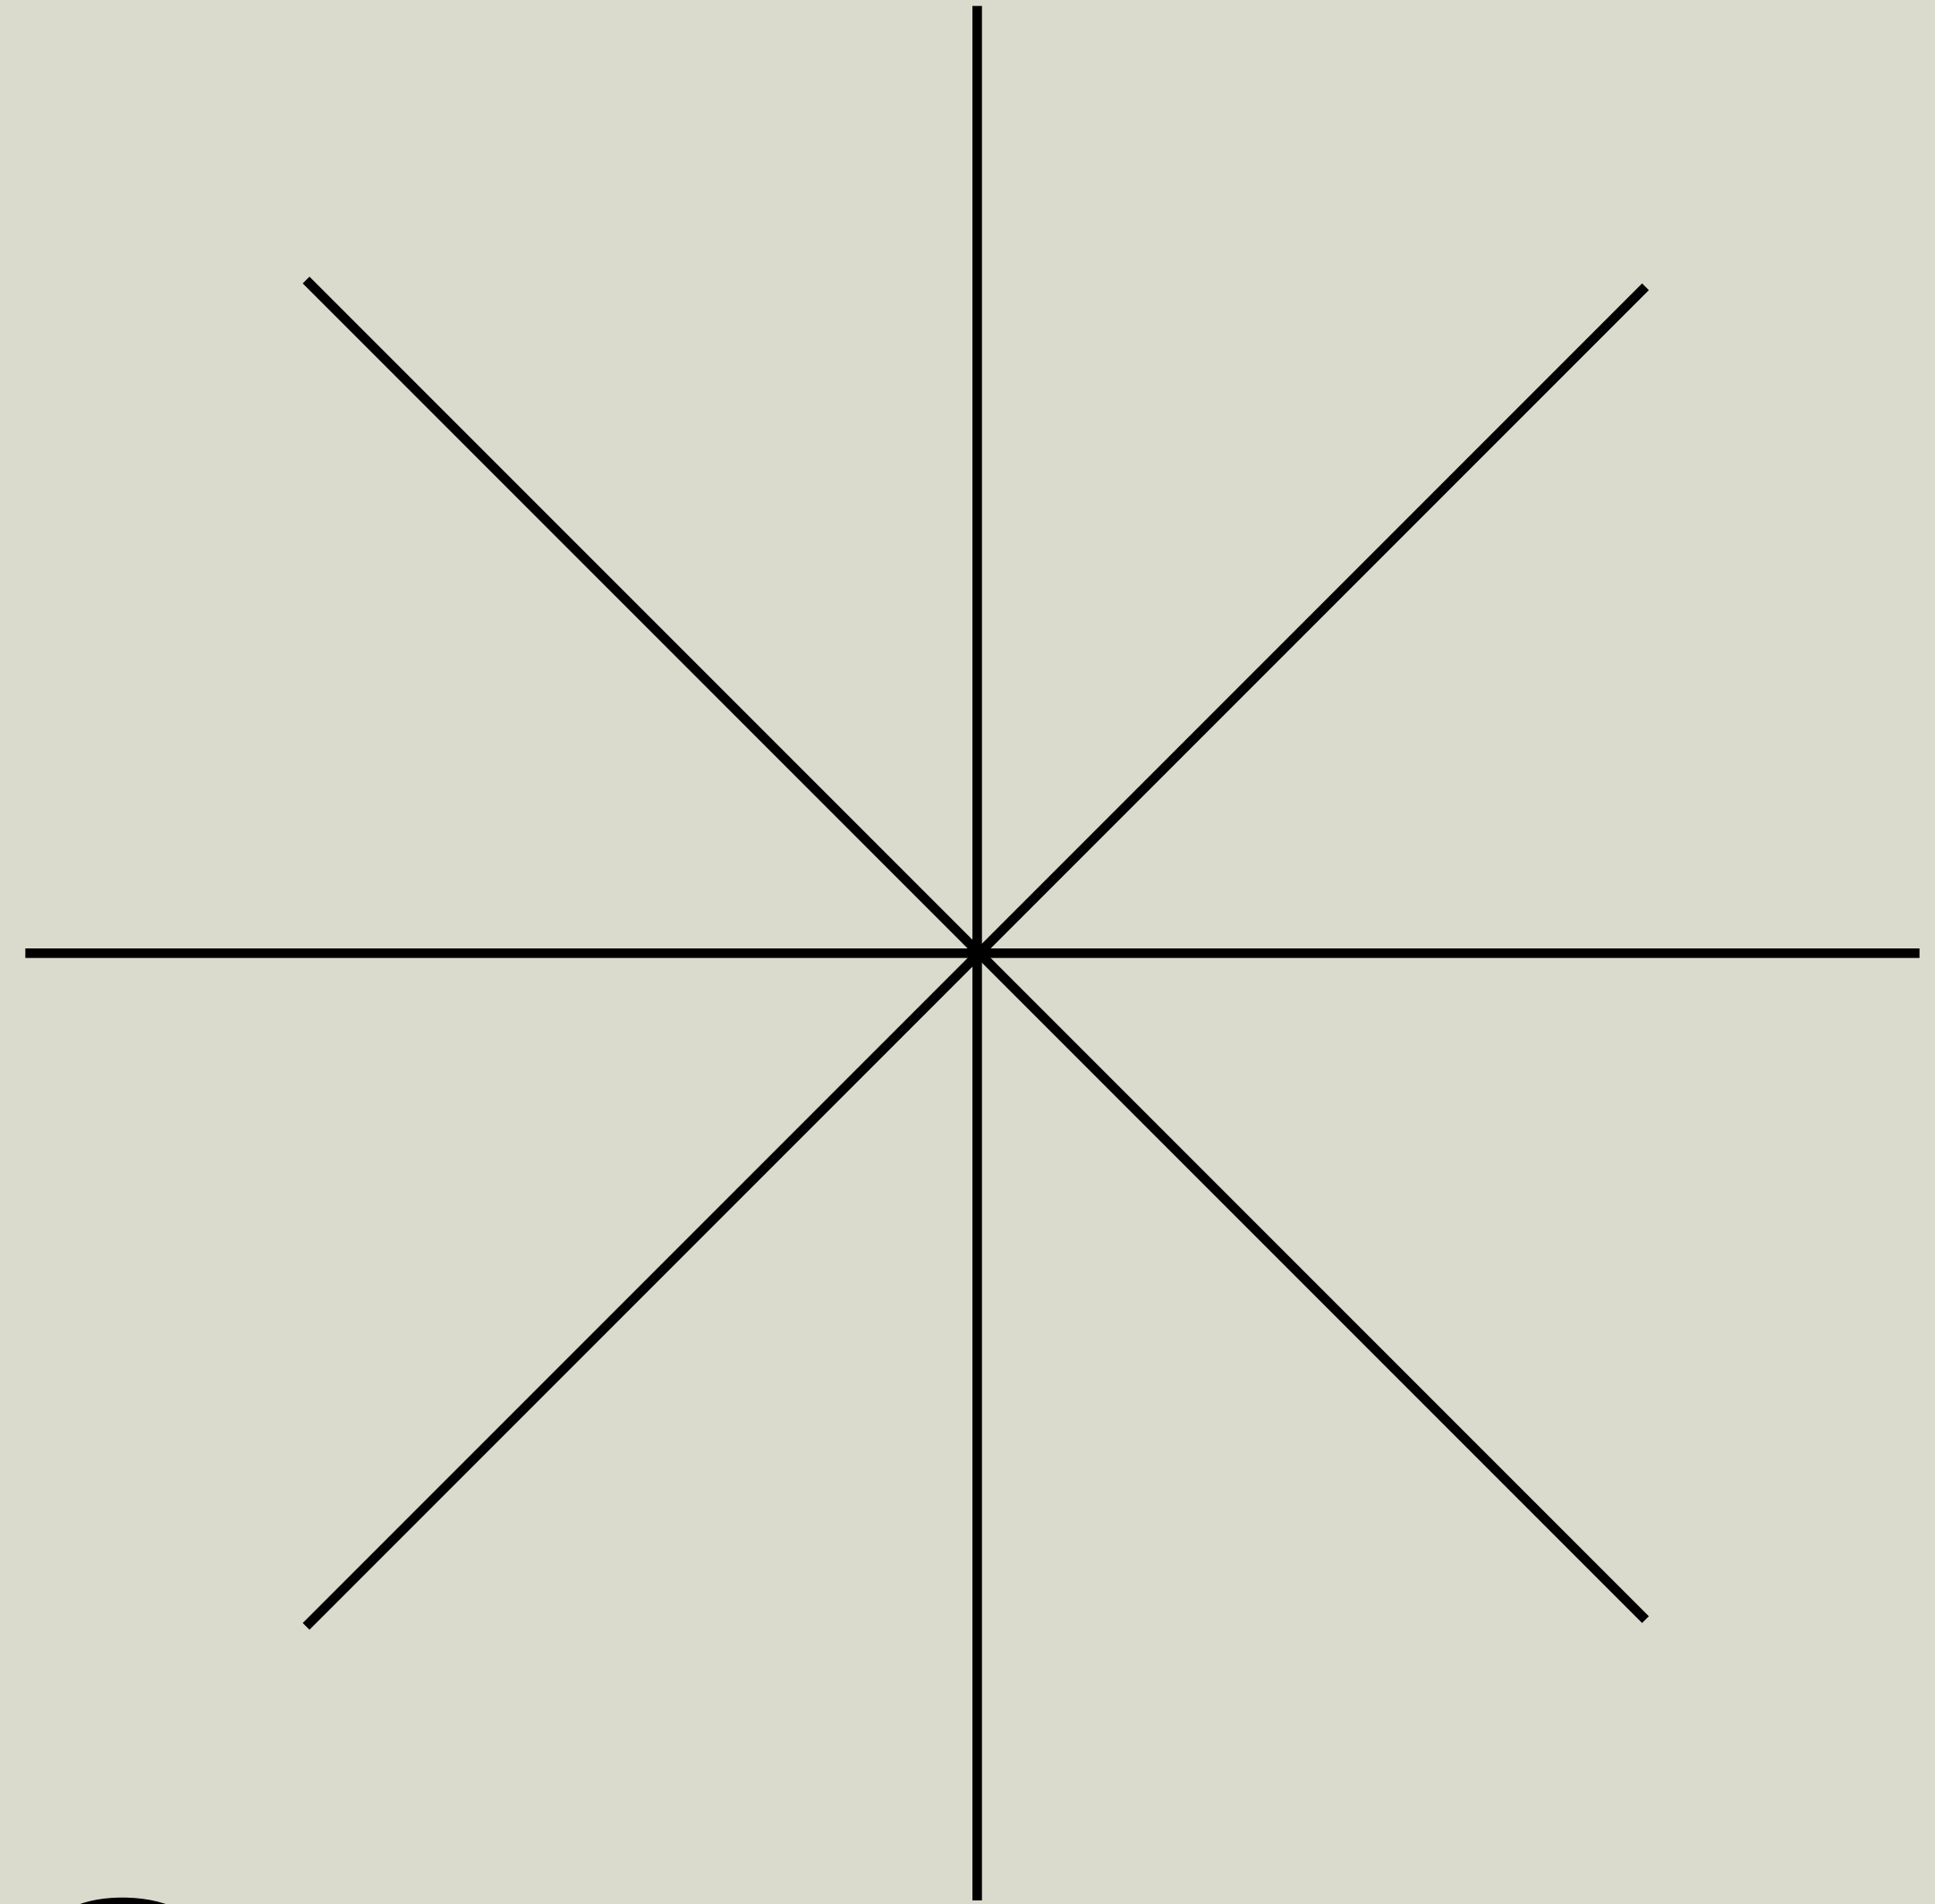
Users can regenerate individual cards or all at once. Locked sentences are preserved. The composition view assembles locked fragments and can polish them into a final paragraph.



From prompt to composition

03

KEY DECISIONS



Design Strategy

PROMPTING

STYLE-DRIVEN PROMPT STRUCTURE

Each card receives the same task, dimensions, and locked sentences, but a different style tag (Direct, Narrative, etc.) This produces meaningful variation while keeping outputs comparable.

CONTROL

VERBATIM LOCK CONSTRAINT

Locked sentences are passed back as Include verbatim: instructions. This constraint preserves user intent without over-specifying the output.

VARIABILITY

NON-DETERMINISM AS A FEATURE

The system treats variability as useful signal, not noise. Each regeneration is an opportunity for the user to find something better. The interface is designed to make comparison easy rather than hide the variance.

UX

DIMENSIONS FOR GUIDED GENERATION

Six sliders (Tone, Length, Creativity, Depth, Persuasion, Emotion) let users shift the generation parameters without writing new prompts, lowering the barrier to iteration.

Initial one-shot prompt:

I would like to build a canvas-based workspace where users can generate multiple approaches to a text simultaneously, compare alternatives side by side, lock elements of the text that they like, and ask the AI to vary specific dimensions, rather than the linear, one-at-a-time interaction model that current chat interfaces impose. It would also be very important to me to be able to see iteration summary and history that has been generated by a user.

Model Used: Claude Sonnet 4.6

⏪
⏩
Paralleltext
canvas

Write an opening paragraph about the future of focused work

REGENERATE ALL
+ 2 MORE
→

Variation Dimensions

Tone Formal · Casual

—————●—————

Length Concise · Verbose

—————●—————

Creativity Literal · Creative

—————●—————

GENERATE WITH THESE SETTINGS

+ VARIANT
VIEW COMPOSITION (4)
CLEAR CANVAS

#01
Direct
50w
✕

The challenge of modern work isn't the volume of tasks—it's the constant interruption of deep focus. Studies show that after a distraction, it takes an average of 23 minutes to return to the original task. This isn't a willpower problem. It's a design problem, and it demands a design solution.

fit82%

#02
Narrative
49w
✕

Picture a craftsperson at their workbench—the kind of absorbed silence that surrounds them when their hands know exactly what to do. That state, psychologists call it flow, is what most of us are chasing through every Pomodoro timer and productivity hack. What if the interface itself could protect it?

fit75%

#03
Analytical
33w
✕

Productivity loss from context-switching represents a measurable cognitive cost: approximately 40% of productive time is lost to task-switching overhead. Existing solutions address symptoms—notification silencing, time-blocking—rather than restructuring the interaction model that generates interruptions.

fit68%

#04
Provocative
42w
✕

We built offices to get people together. Then we built walls to separate them. Then open plans to reunite them. Now we hand everyone a notification firehose and wonder why nothing gets done. The problem was never proximity. It was always attention.

fit91%

Locked Elements

That state, psychologists call it flow, is what most of us are chasing through every Pomodoro timer and productivity hack.

✕

The challenge of modern work isn't the volume of tasks—it's the constant

✕

Vary a Dimension

INSTRUCTION

e.g. "Make it more urgent" or "Shift to second person"

APPLY VARIATION →

Iteration Log

iter 1 01:24 PM

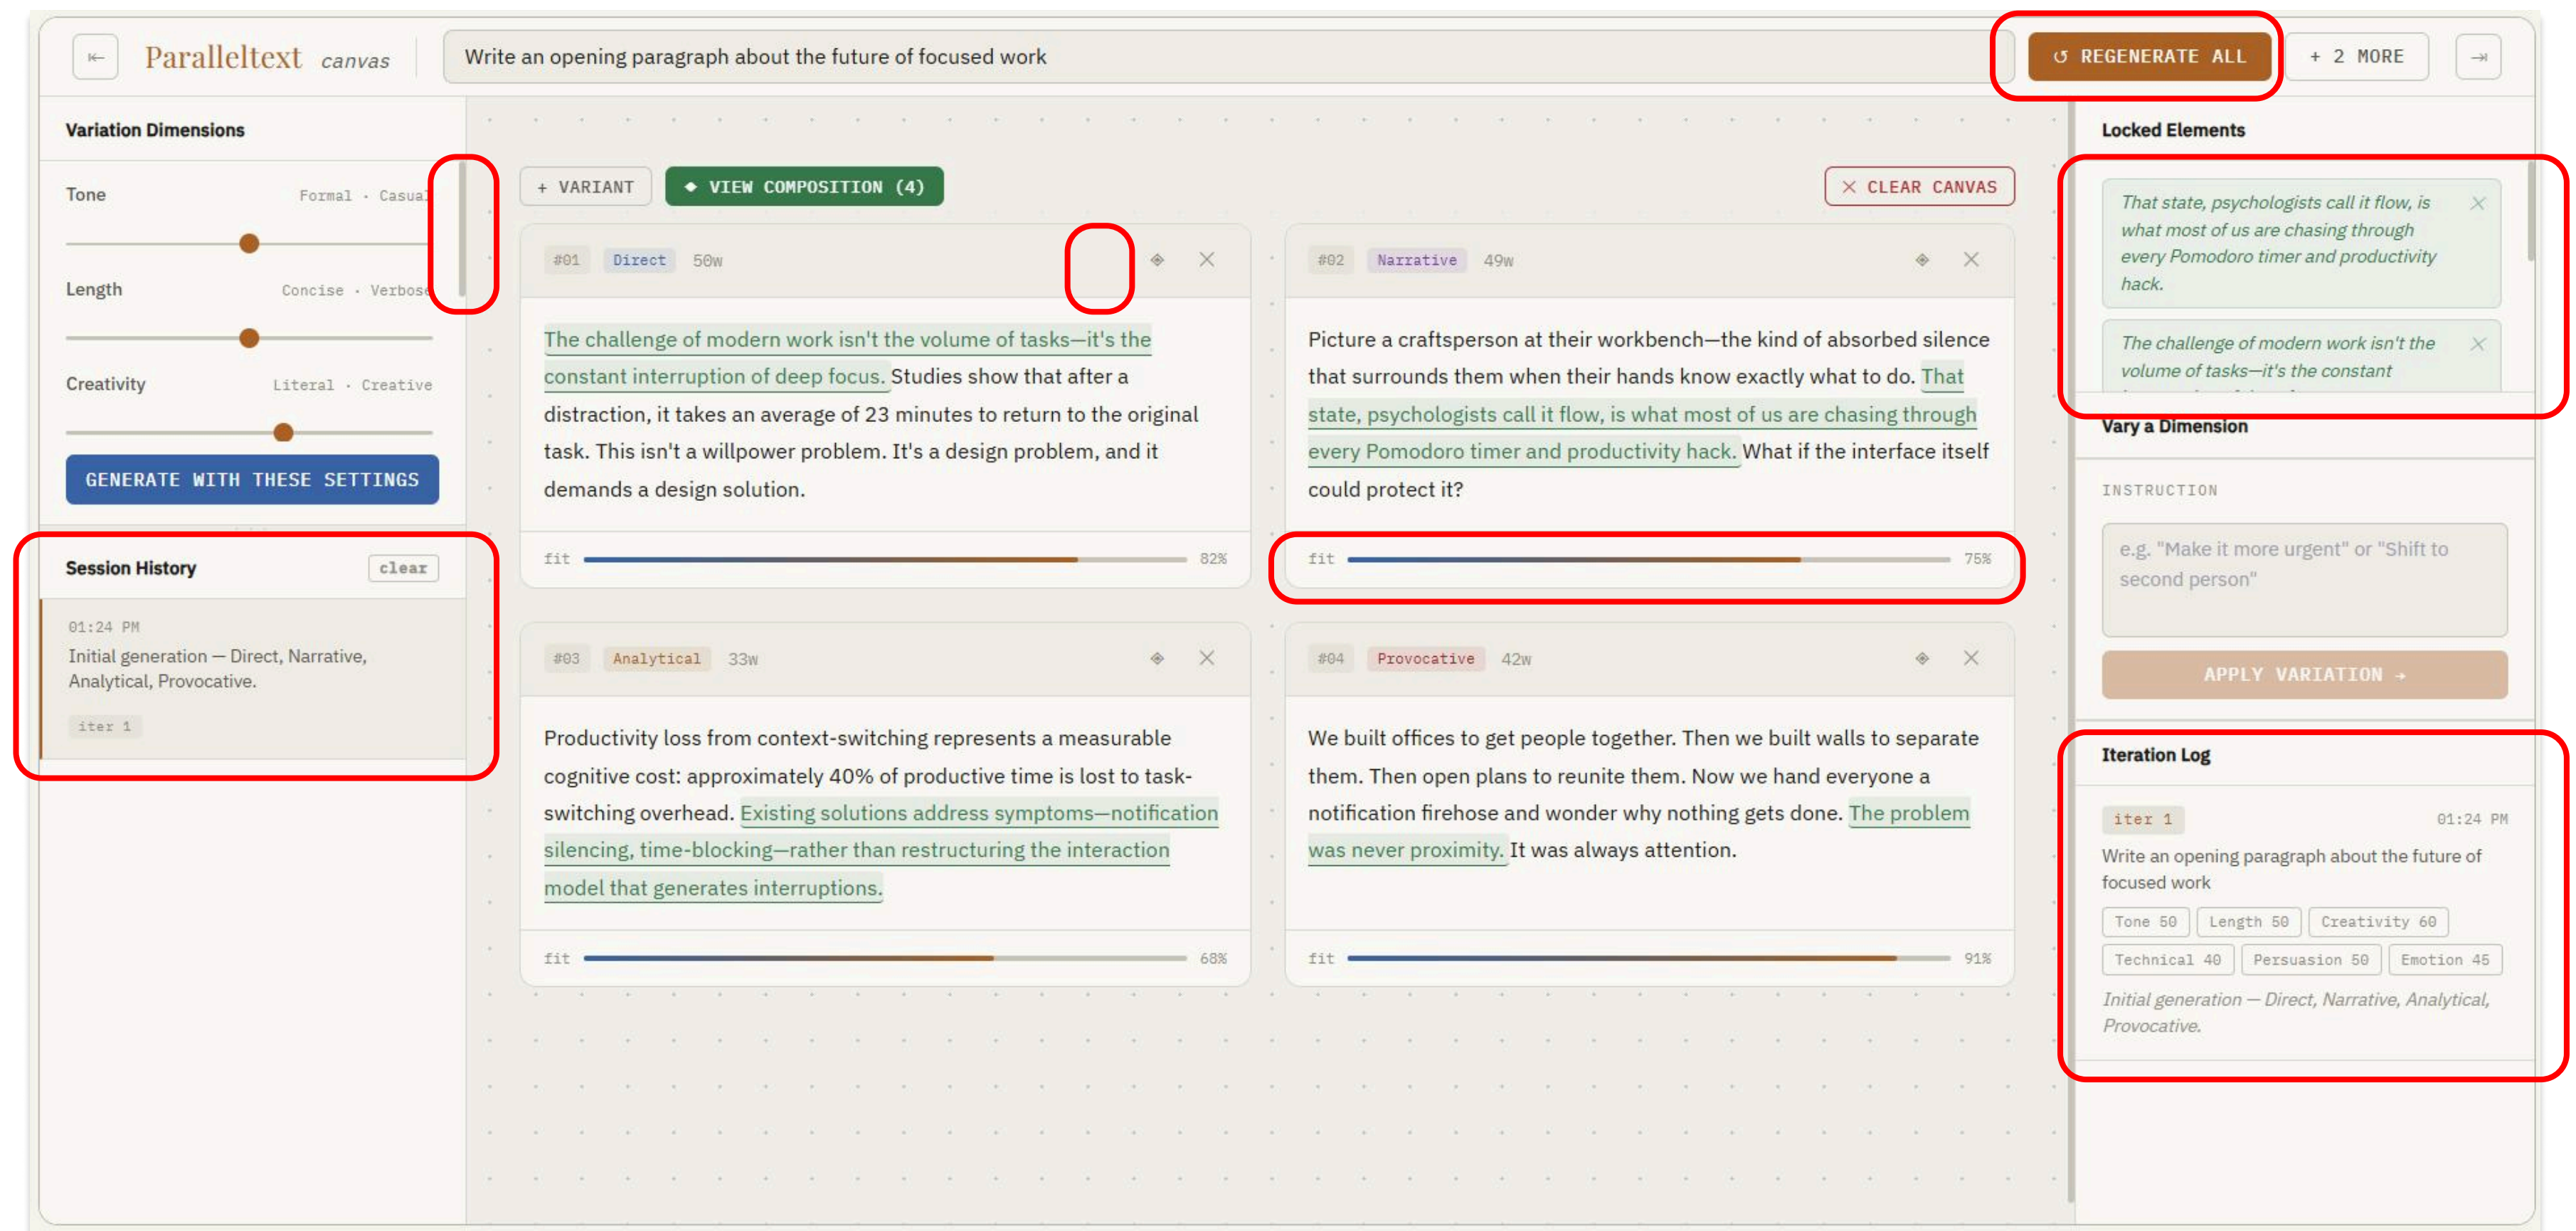
Write an opening paragraph about the future of focused work

Tone 50
Length 50
Creativity 60

Technical 40
Persuasion 50
Emotion 45

Initial generation — Direct, Narrative, Analytical, Provocative.

- No need to separate ‘session history’ from ‘iteration log’
- the variation dimensions and locked elements sections too small, and are not able to cleanly show both their contents and CTAs
- there is no ability to regenerate one card at a time, only the ability to regenerate all cards at once
- even though there was a ‘regenerate all’ button, it did not work when clicked
- the “fit” metric was hard to understand; meaning was unclear



My Prompting Strategy:

goal was to make
focused requests to
ensure main
functionality of the
experience

It was important to me
to scope down not only
the overall project, but
my prompting requests
as well, for the sake of
feasibility

I also need to make sure that when I hit the 'regenerate all' button, the 'locked' elements actually stay locked, and do not change.

Looking at the code, the issue is that `handleRegenerateAll` calls `runGeneration` with `replacingAll: true`, which wipes all cards and replaces them with fresh placeholders — including any cards whose sentences were locked. The fix is to keep locked cards in place and only replace the unlocked ones.

text-canvas-workspace.jsx
Interactive artifact • Version 6

Now when you hit **Regenerate All**:

- Any card contain[↓] a locked sentence is preserved exactly as-is

Reply...

+ Sonnet 4.6 ▾

The screenshot shows a writing assistant interface with a central canvas and several side panels. The main canvas displays a composition of four text blocks. The left panel, titled 'DIMENSIONS', contains sliders for Tone, Length, Creativity, Depth, Persuasion, and Emotion, along with a 'Generate with settings' button. The right panel, titled 'LOCKED', shows a list of generated text snippets with 'X' icons for removal. The bottom-left panel, titled 'LOG', shows the generation process and a list of approaches (T50, L50, C60, D40, P50, E45). The bottom-right panel, titled 'VARY', shows a text input field and an 'Apply' button. Red annotations highlight the 'Generate with settings' button, the 'DIMENSIONS' panel, the 'LOG' panel, the 'LOCKED' panel, and the 'VARY' panel. Additionally, a red circle highlights a refresh icon in the top-right corner of the first text block, and a red rectangle highlights the bottom of the 'Analytical' text block.

← Paths canvas | Write an opening paragraph about the future of focused work [Regenerate] + 2 more →

DIMENSIONS

Tone: Formal / Casual

Length: Concise / Verbose

Creativity: Literal / Creative

Depth: Simple / Technical

Persuasion: Neutral / Persuasive

Emotion: Reserved / Emotive

Generate with settings

Composition (4) [Clear]

Direct 50w [Refresh] [Close]

The challenge of modern work isn't the volume of tasks—it's the constant interruption of deep focus. Studies show that after a distraction, it takes an average of 23 minutes to return to the original task. This isn't a willpower problem. It's a design problem, and it demands a design solution.

Narrative 49w [Refresh] [Close]

Picture a craftsperson at their workbench—the kind of absorbed silence that surrounds them when their hands know exactly what to do. That state, psychologists call it flow, is what most of us are chasing through every Pomodoro timer and productivity hack. What if the interface itself could protect it?

Analytical 32w [Refresh] [Close]

Productivity loss from context-switching represents a measurable cognitive cost: approximately 40% of productive time is lost to task-switching overhead. Existing solutions address symptoms rather than restructuring the interaction model that generates interruptions.

Provocative 42w [Refresh] [Close]

We built offices to get people together. Then we built walls to separate them. Then open plans to reunite them. Now we hand everyone a notification firehose and wonder why nothing gets done. The problem was never proximity. It was always attention.

LOCKED

Studies show that after a distraction, it takes an average of 23 minutes to return to the original task. [Close]

Then we built walls to separate them. [Close]

Existing solutions address symptoms rather than restructuring the interaction model that generates interruptions. [Close]

Picture a craftsperson at their workbench—the kind of absorbed silence that surrounds them when their hands know exactly what to do. [Close]

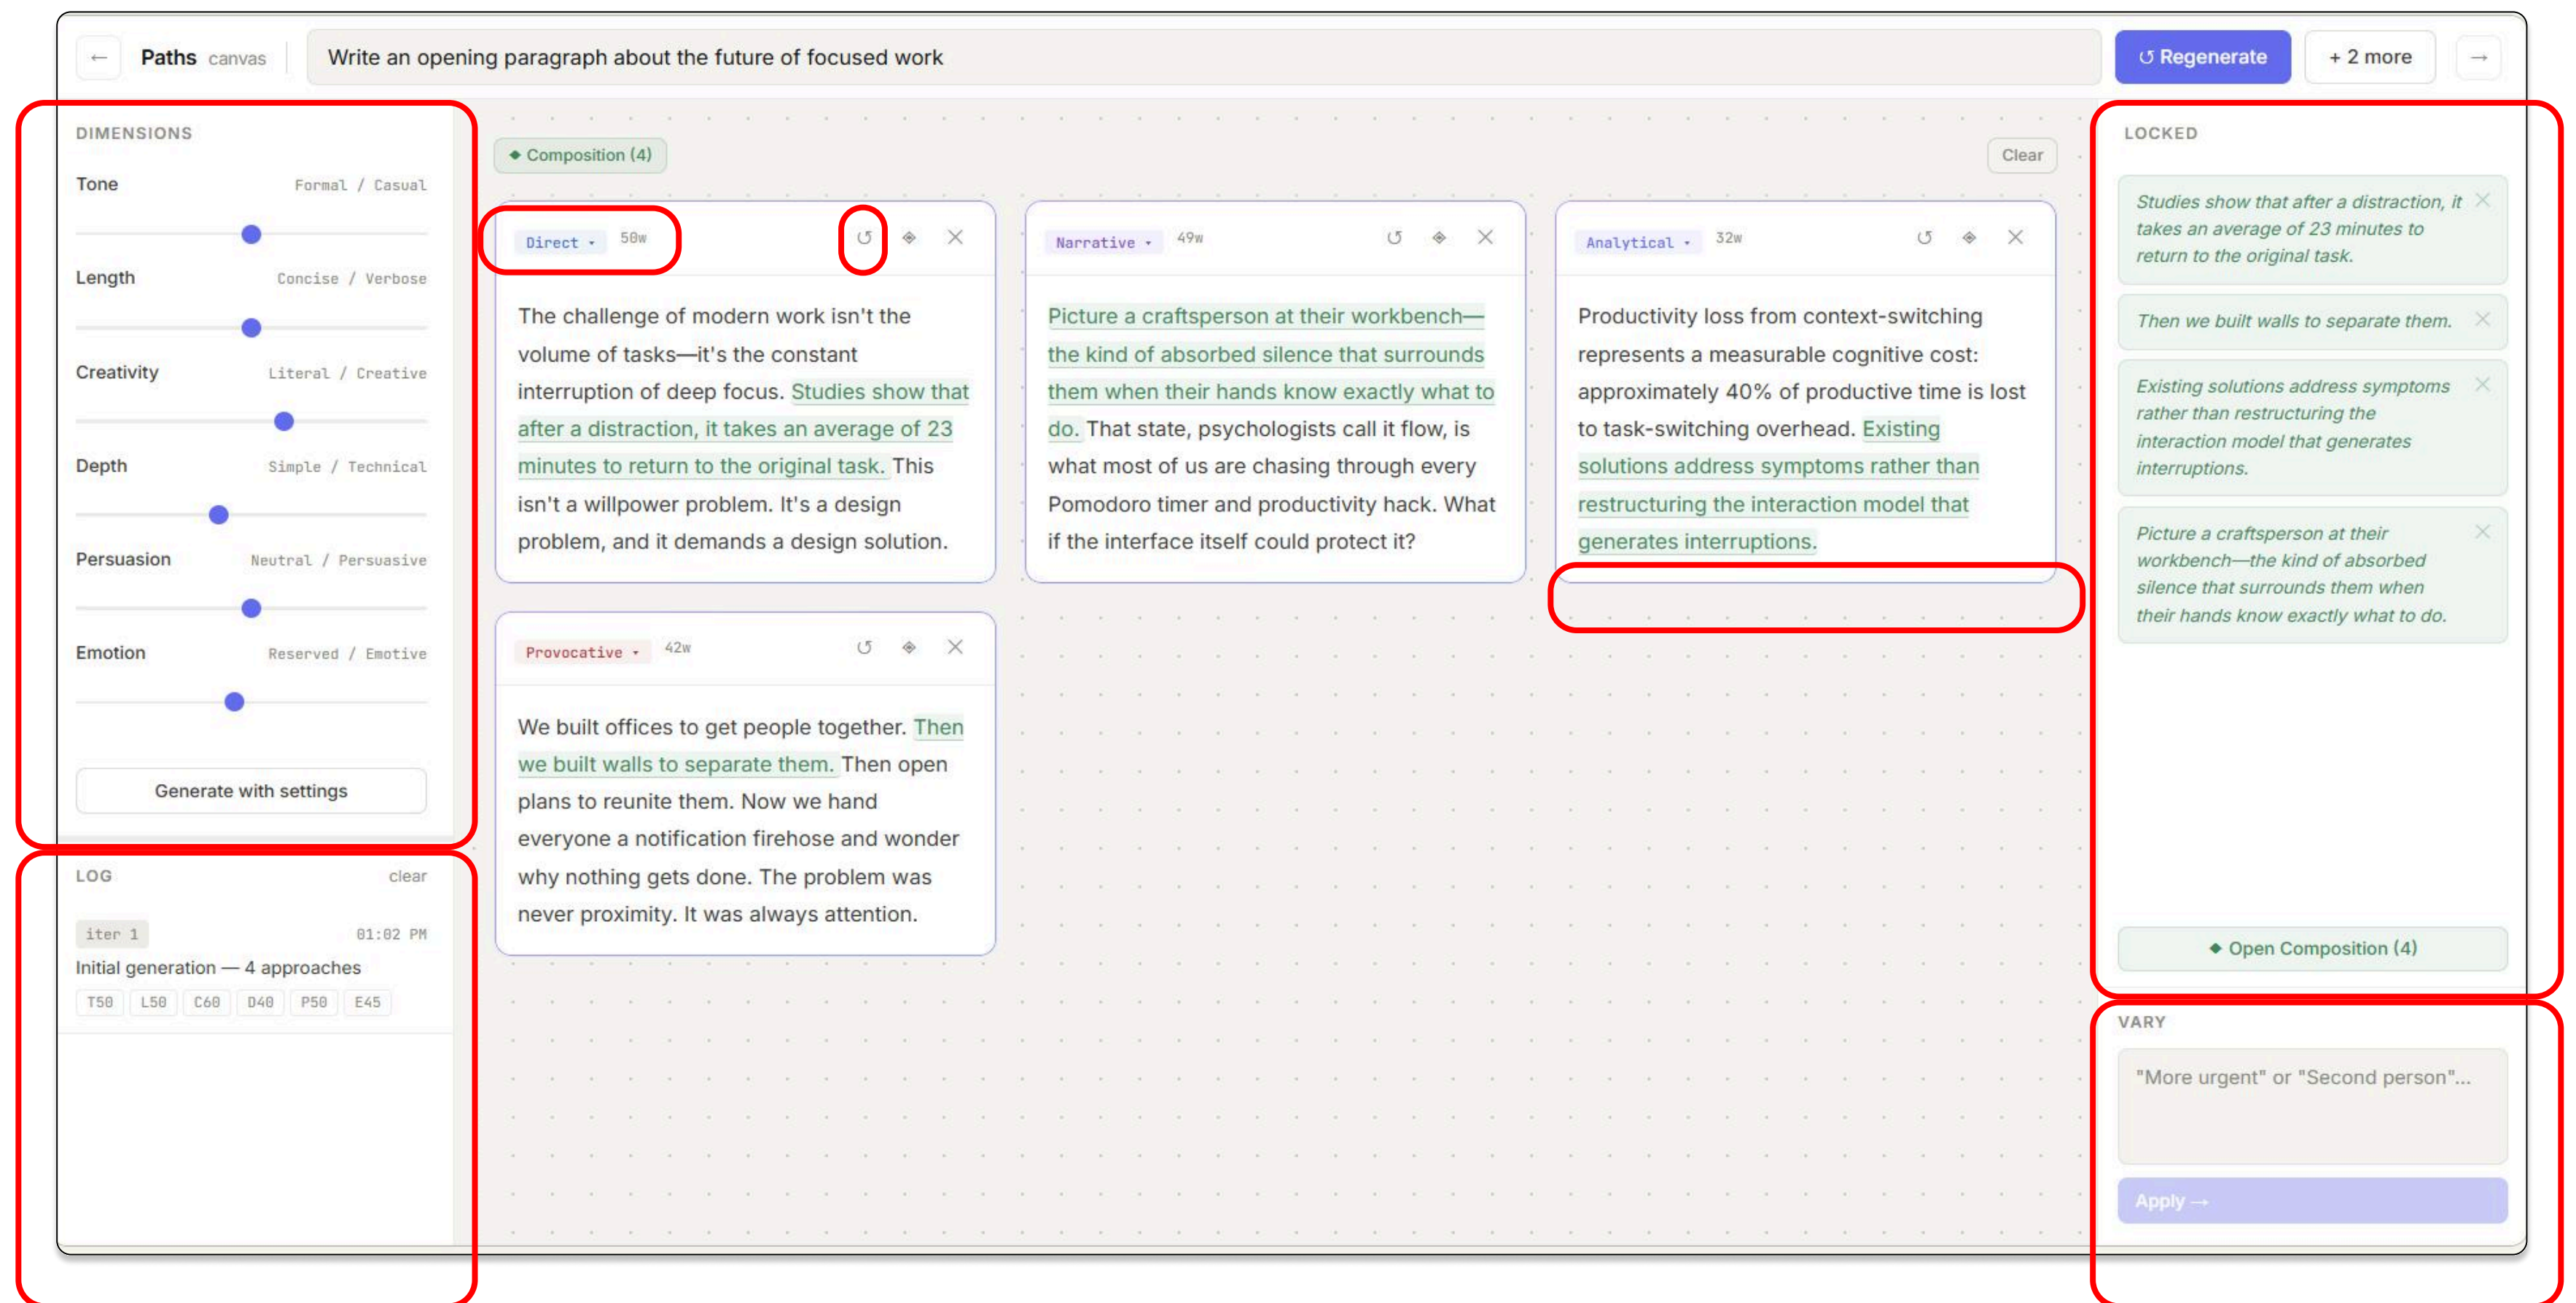
◆ Open Composition (4)

VARY

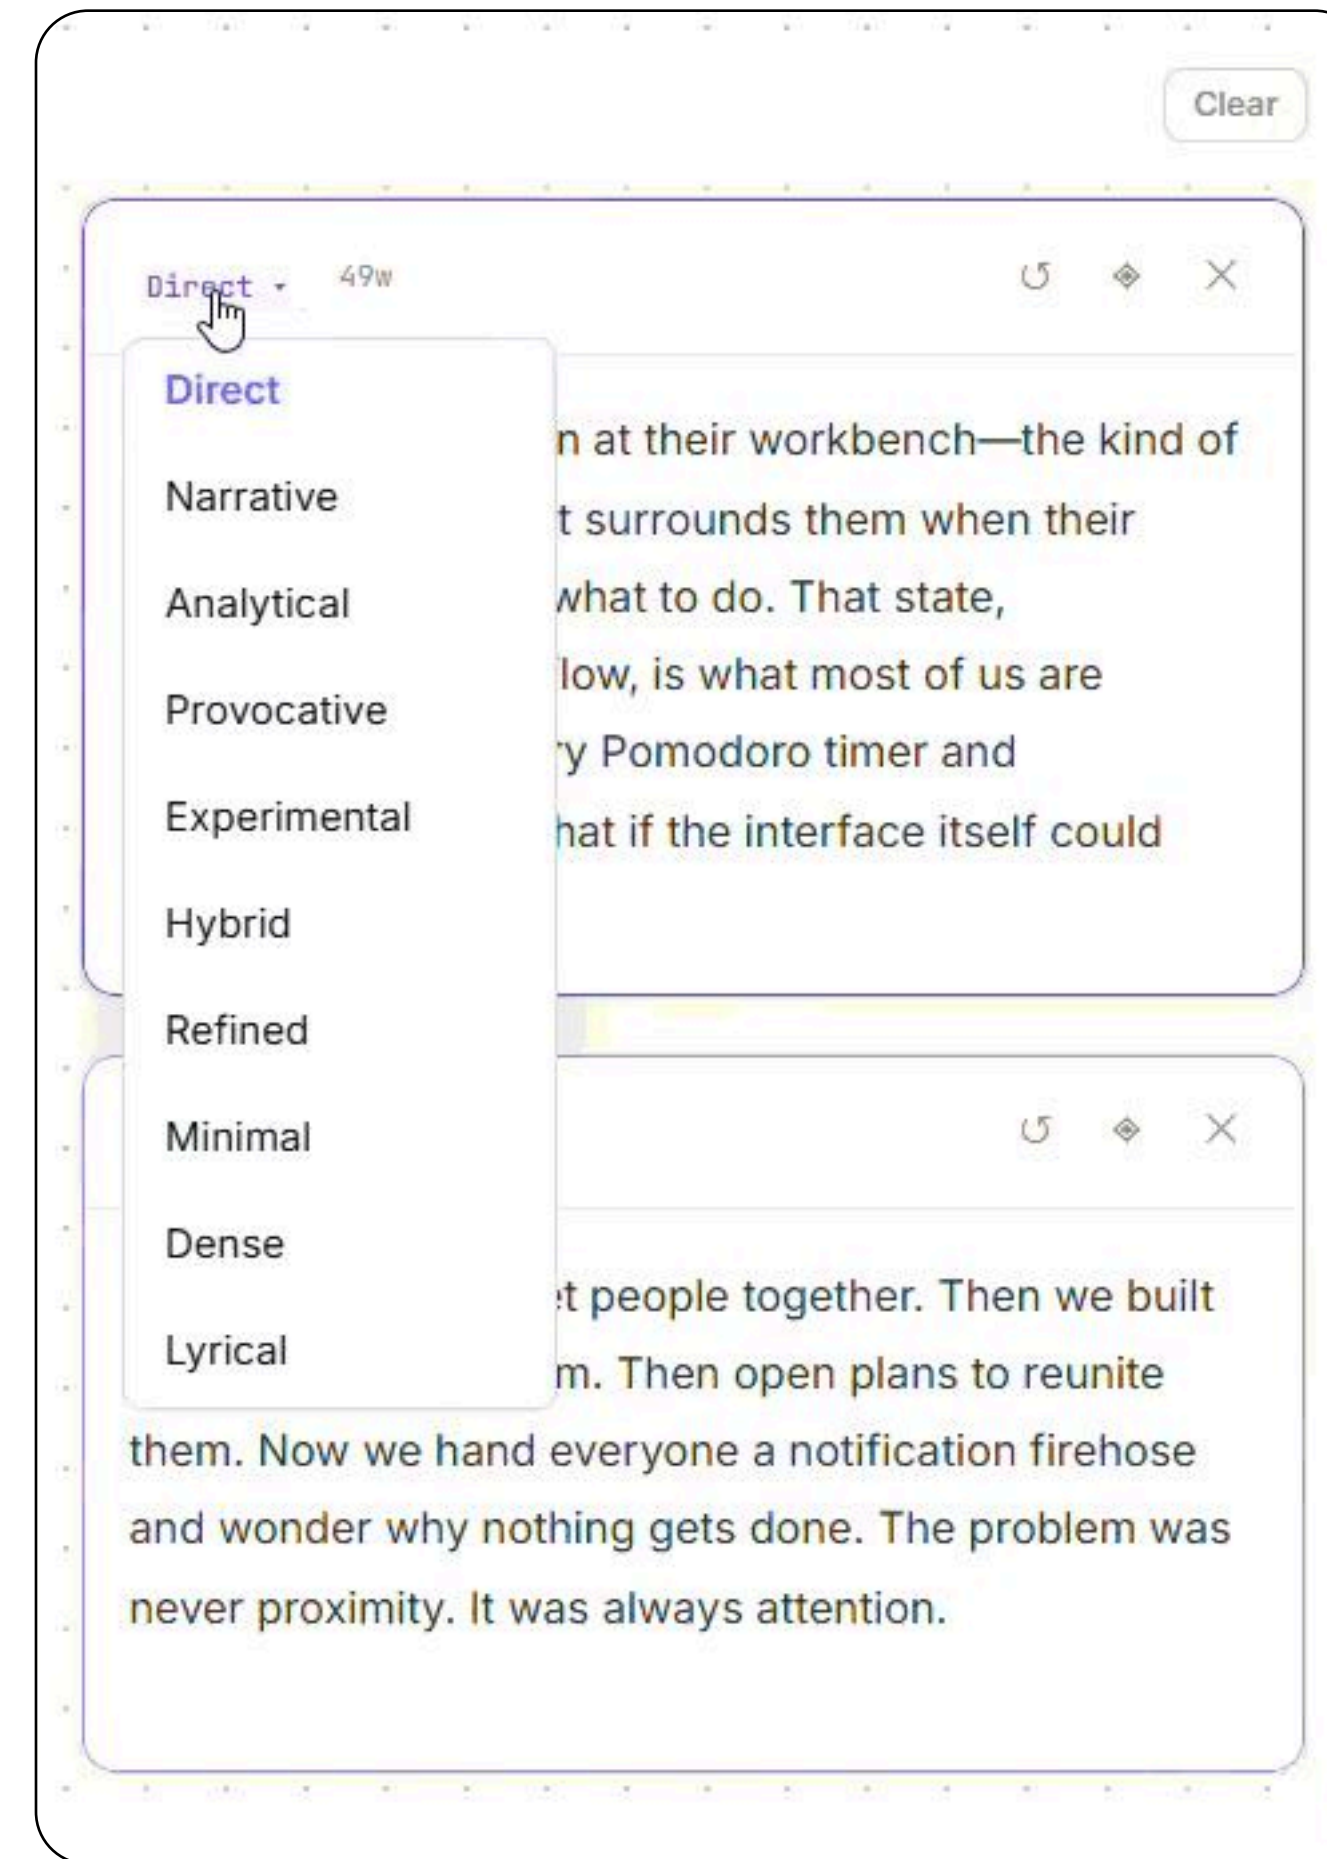
"More urgent" or "Second person"...

Apply →

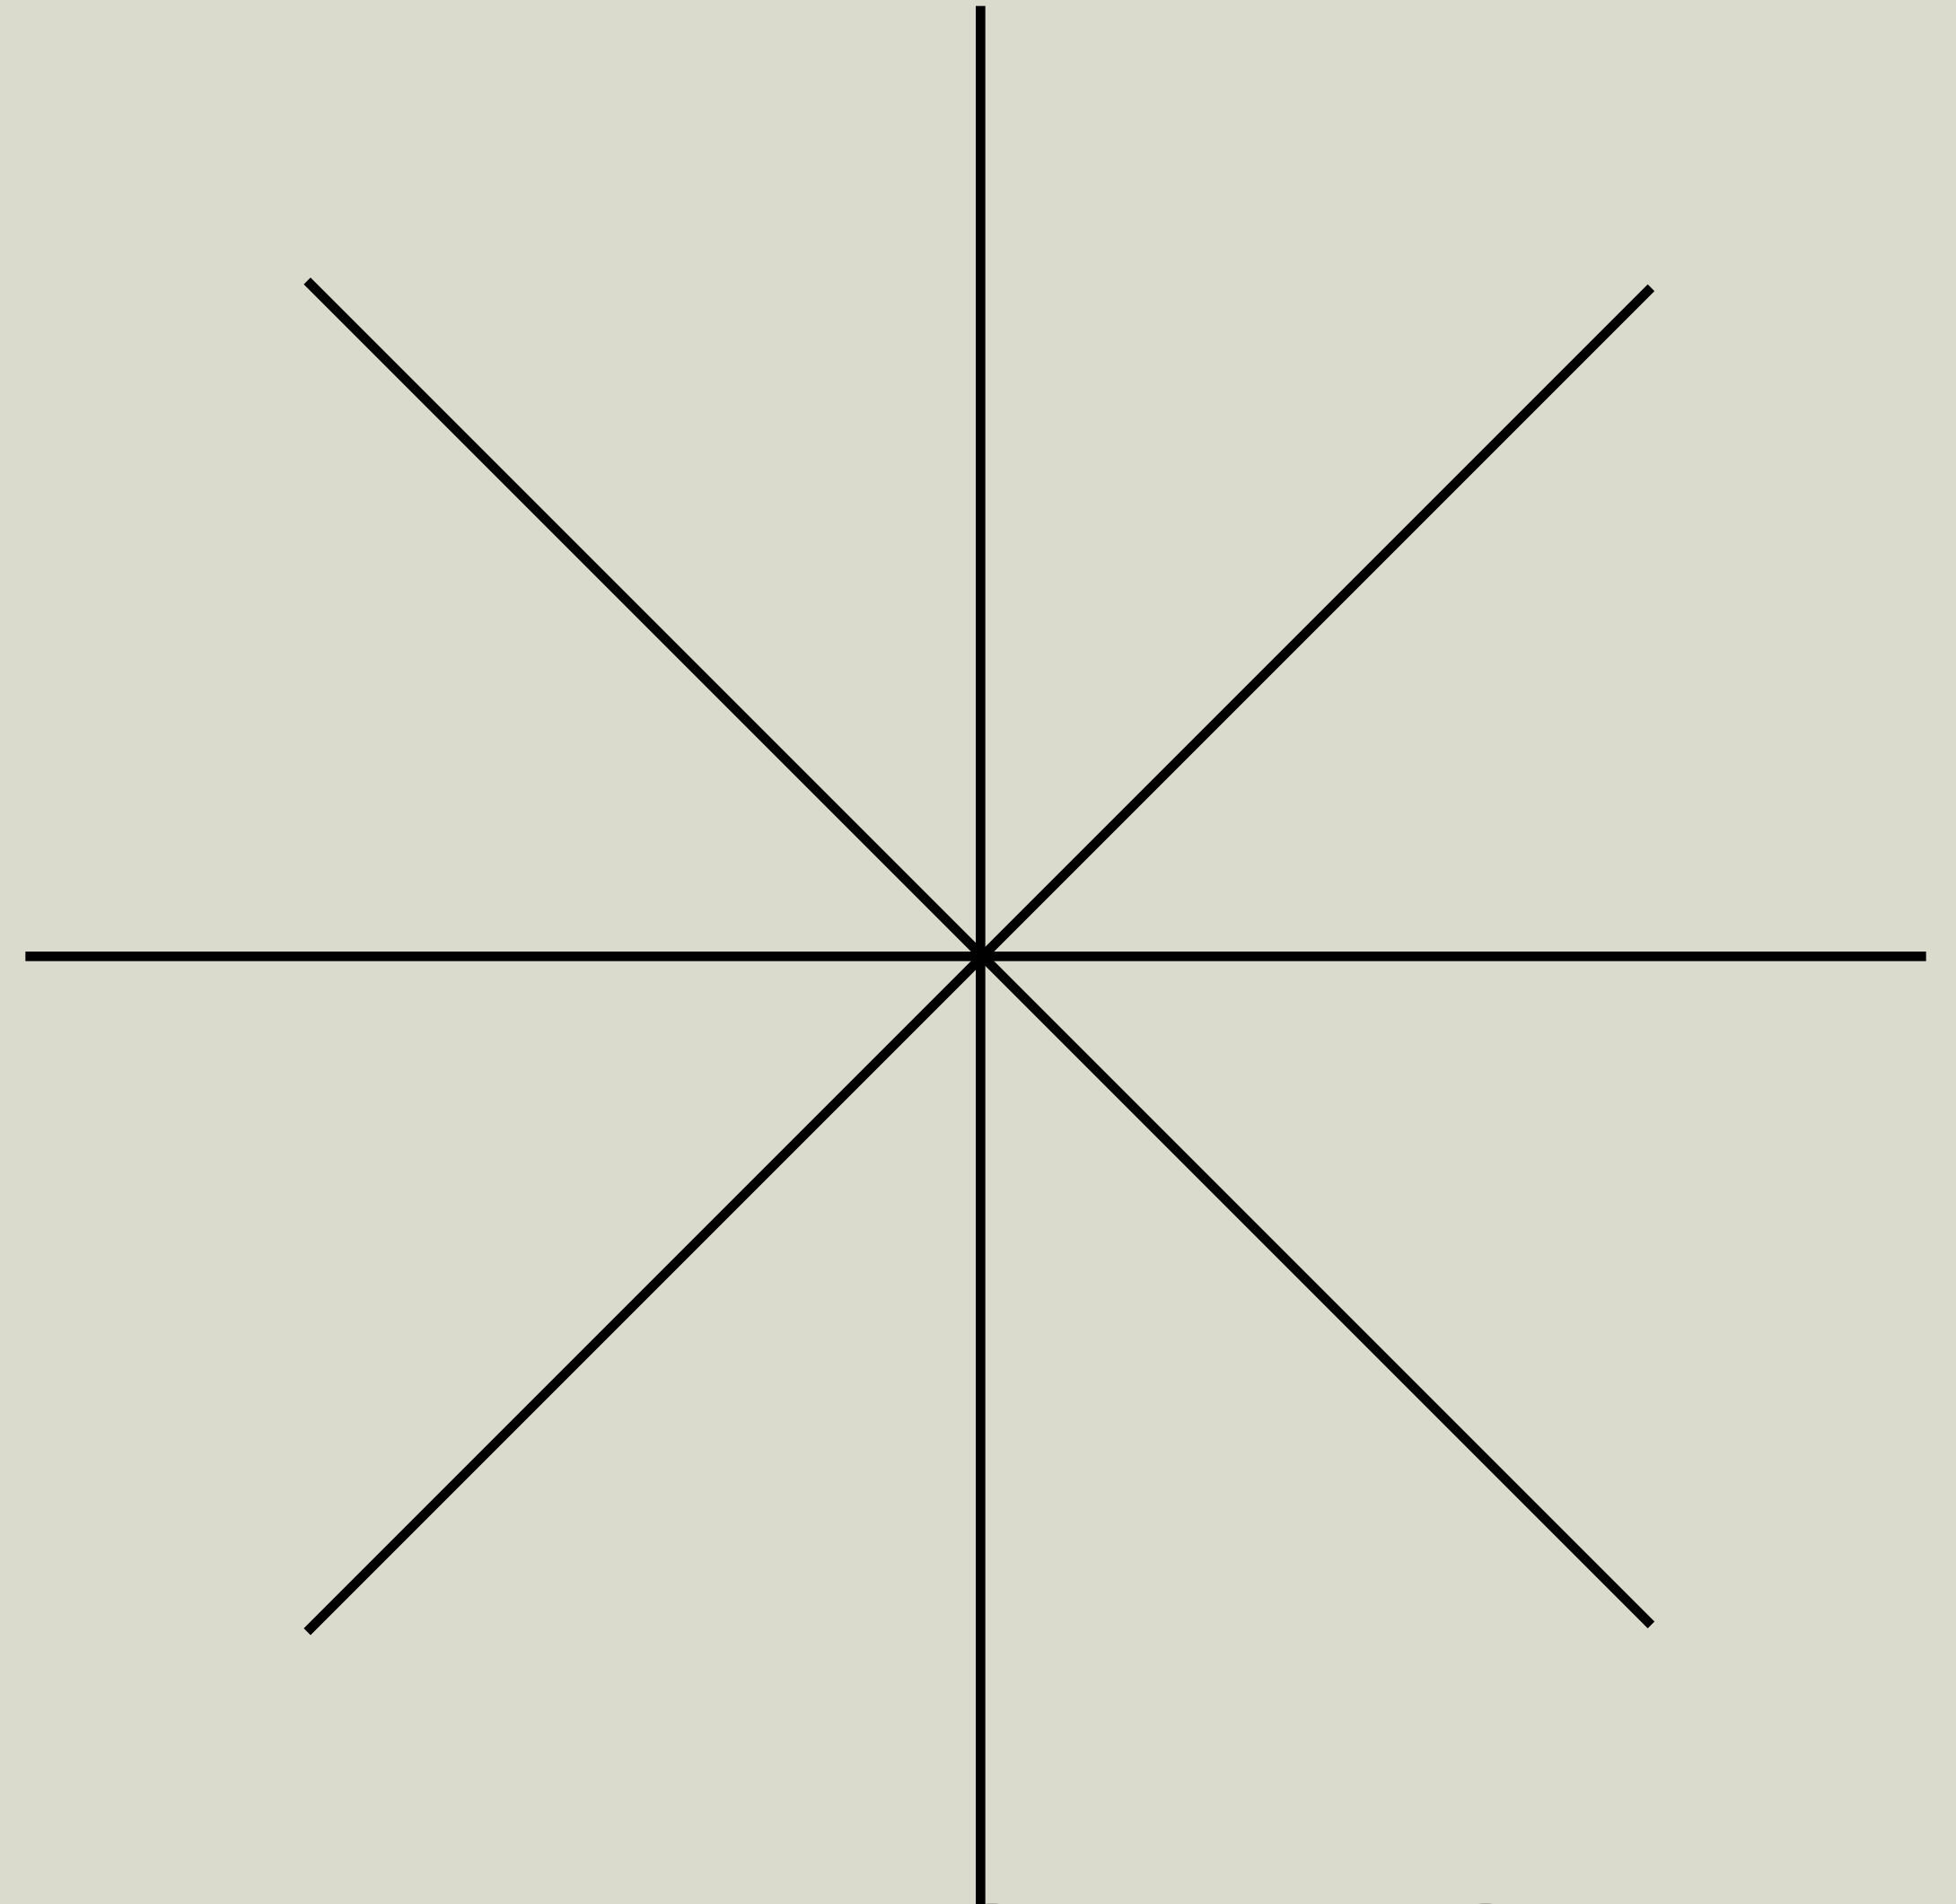
- Still using Claude
- Several small prompting tweaks to get final result
- Quick usability testing was helpful for cleaning up small UI issues, such as the need to scroll to see all the dimensions and locked elements.
- Eval testing helped to ensure main functionality was working as desired.
- Removed 'fit' metric as its meaning was unclear
- removed duplicate session logs
- ensured users could regenerate one card at a time
- ensured users could see CTAs in properties panels



- Worked on developing smaller features, such as the ability to change the style of one card at a time, as well as to lock something and then regenerate one card at a time



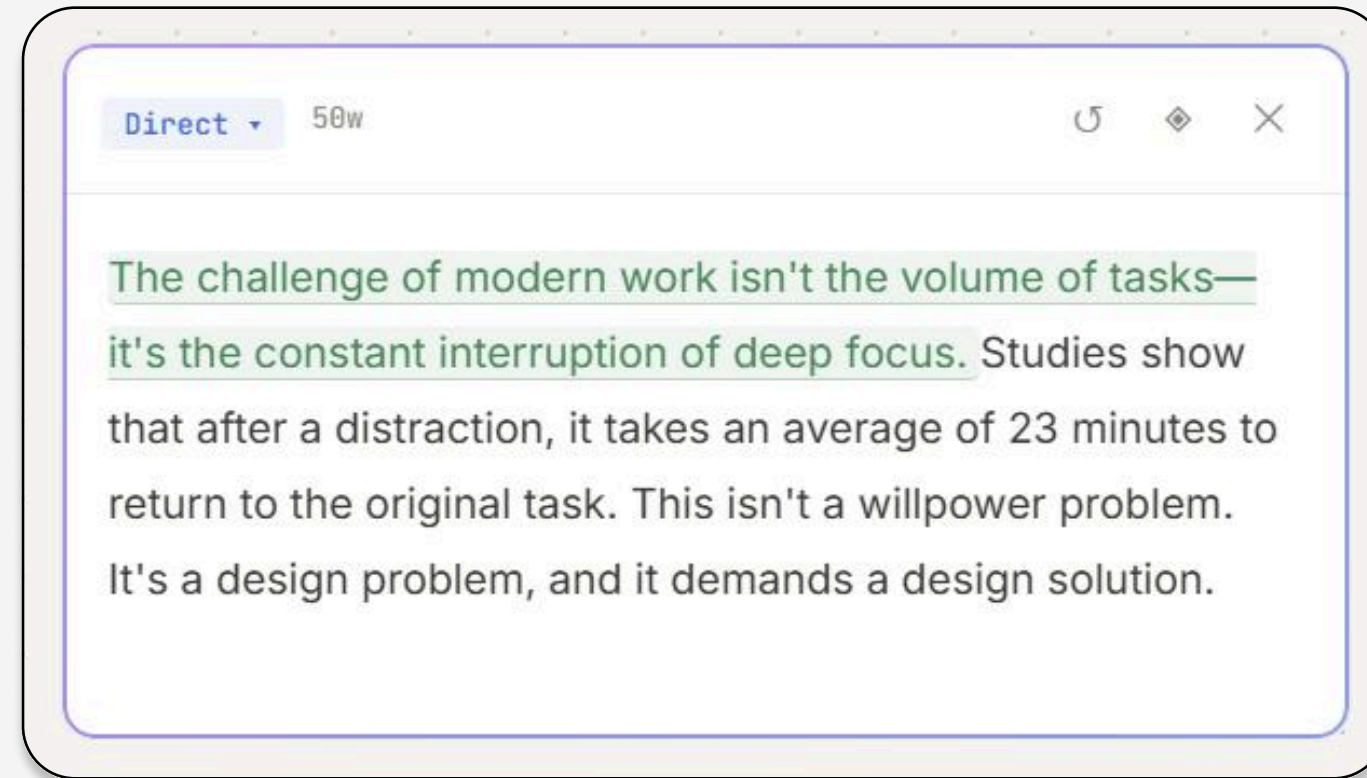
developing more user control over a single card



04

OBSERVED BEHAVIORS

Eval Testing



When a user locks a sentence and regenerates a card, does the model **preserve** that sentence verbatim?

METRIC 1 - PRESERVATION

Exact Match

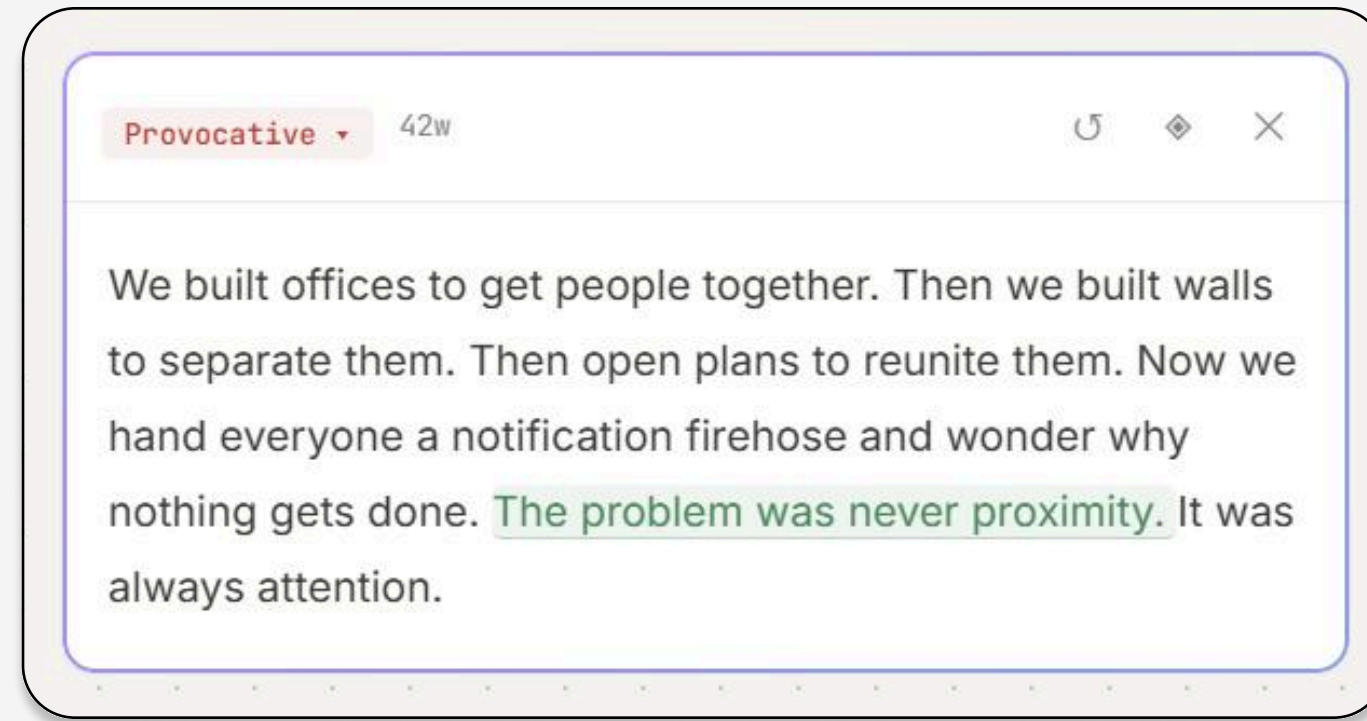
Does the locked sentence appear character-for-character in the output? Checked programmatically. Returns yes / no per run.

METRIC 2 - COHERENCE

LLM-as-Judge

Does the surrounding text flow naturally around the lock? Scored 1-5 by a second Claude call using a rubric anchored to seam quality.

How the eval was run



When a user locks a sentence and regenerates a card, does the model **preserve** that sentence verbatim?

[LINK TO EVAL PROTOTYPE](#)

CONFIGURATION

Prompt	Write an opening paragraph about the future of focused work.
Style Tag	Provocative
Locked Sentence	“The problem was never proximity. It was always attention.”
Runs	15 regenerations
System Prompt	Exact Paths source (SYSTEM constant)
User Prompt	Exact buildUserPrompt() from Paths source

PER-RUN FLOW

- 1 - Generate**
Send prompt + lock to Claude using exact Paths prompts.
- 2 - Check Preservation**
Exact string match - is the locked sentence in the output?
- 3 - Judge Coherence**
Second Claude call scores 1-5 with one-sentence rationale.
- 4 - If Failed**
Coherence marked N/A - preserves integrity of the metric.

15 RUNS COMPLETED

Paths Canvas
LOCKED ELEMENT EVAL · SCENARIO 2 · 15 RUNS

Write an opening paragraph about the future of focused work.

STYLE TAG
Direct Narrative Analytical Provocative Experimental

The locked sentence will only be injected into regenerations of this style card - matching Paths behaviour.

LOCKED SENTENCE
The problem was never proximity. It was always attention.
Passed to the model exactly as: Include verbatim: "[sentence]"

PROGRESS: 15 / 15
PRESERVED: 14
FAILED: 1
AVG COHERENCE: 5.0

RUN	GENERATED OUTPUT	PRESERVED?	COHERENCE
01	The future of focused work won't be determined by the tools we build or the offices we abandon, but by whether we can remember what concentration actually feels like. <i>The problem was never proximity. It was always attention.</i> We've spent decades optimizing for collaboration while forgetting that the rarest resource isn't connection—it's the ability to think alone, uninterrupted, for more than seven minutes at a time.	✓ yes	5/5 <i>The locked sentence serves as a perfect hinge between the opening claim and the concluding observation, with both surrounding sentences flowing naturally into and out of it as if the entire paragraph were written as one cohesive piece.</i>

[LINK TO EVAL PROTOTYPE](#)

INTEPRETATION

The include verbatim: instruction in Paths is highly effective - 93% of regenerations preserved the locked sentence exactly. A 5.0 coherence score across all preserved runs suggested the model is not mechanically inserting the sentence but writing naturally around it.

14/15
Preserved
93%

1
Failed
7%

5.0
Avg. Coherence
out of 5

“The problem was never proximity. It was always attention.”

≠

“The problem was never proximity-it was always attention.”

**Semantic compliance,
not syntactic**

Locked Sentence
(Expected)

What the model produced
(Run 15)

WHAT DOES THIS MEAN?

The only difference was a period and a space being replaced by an em dash.

The meaning is identical. The wording is nearly identical. But the verbatim check fails.

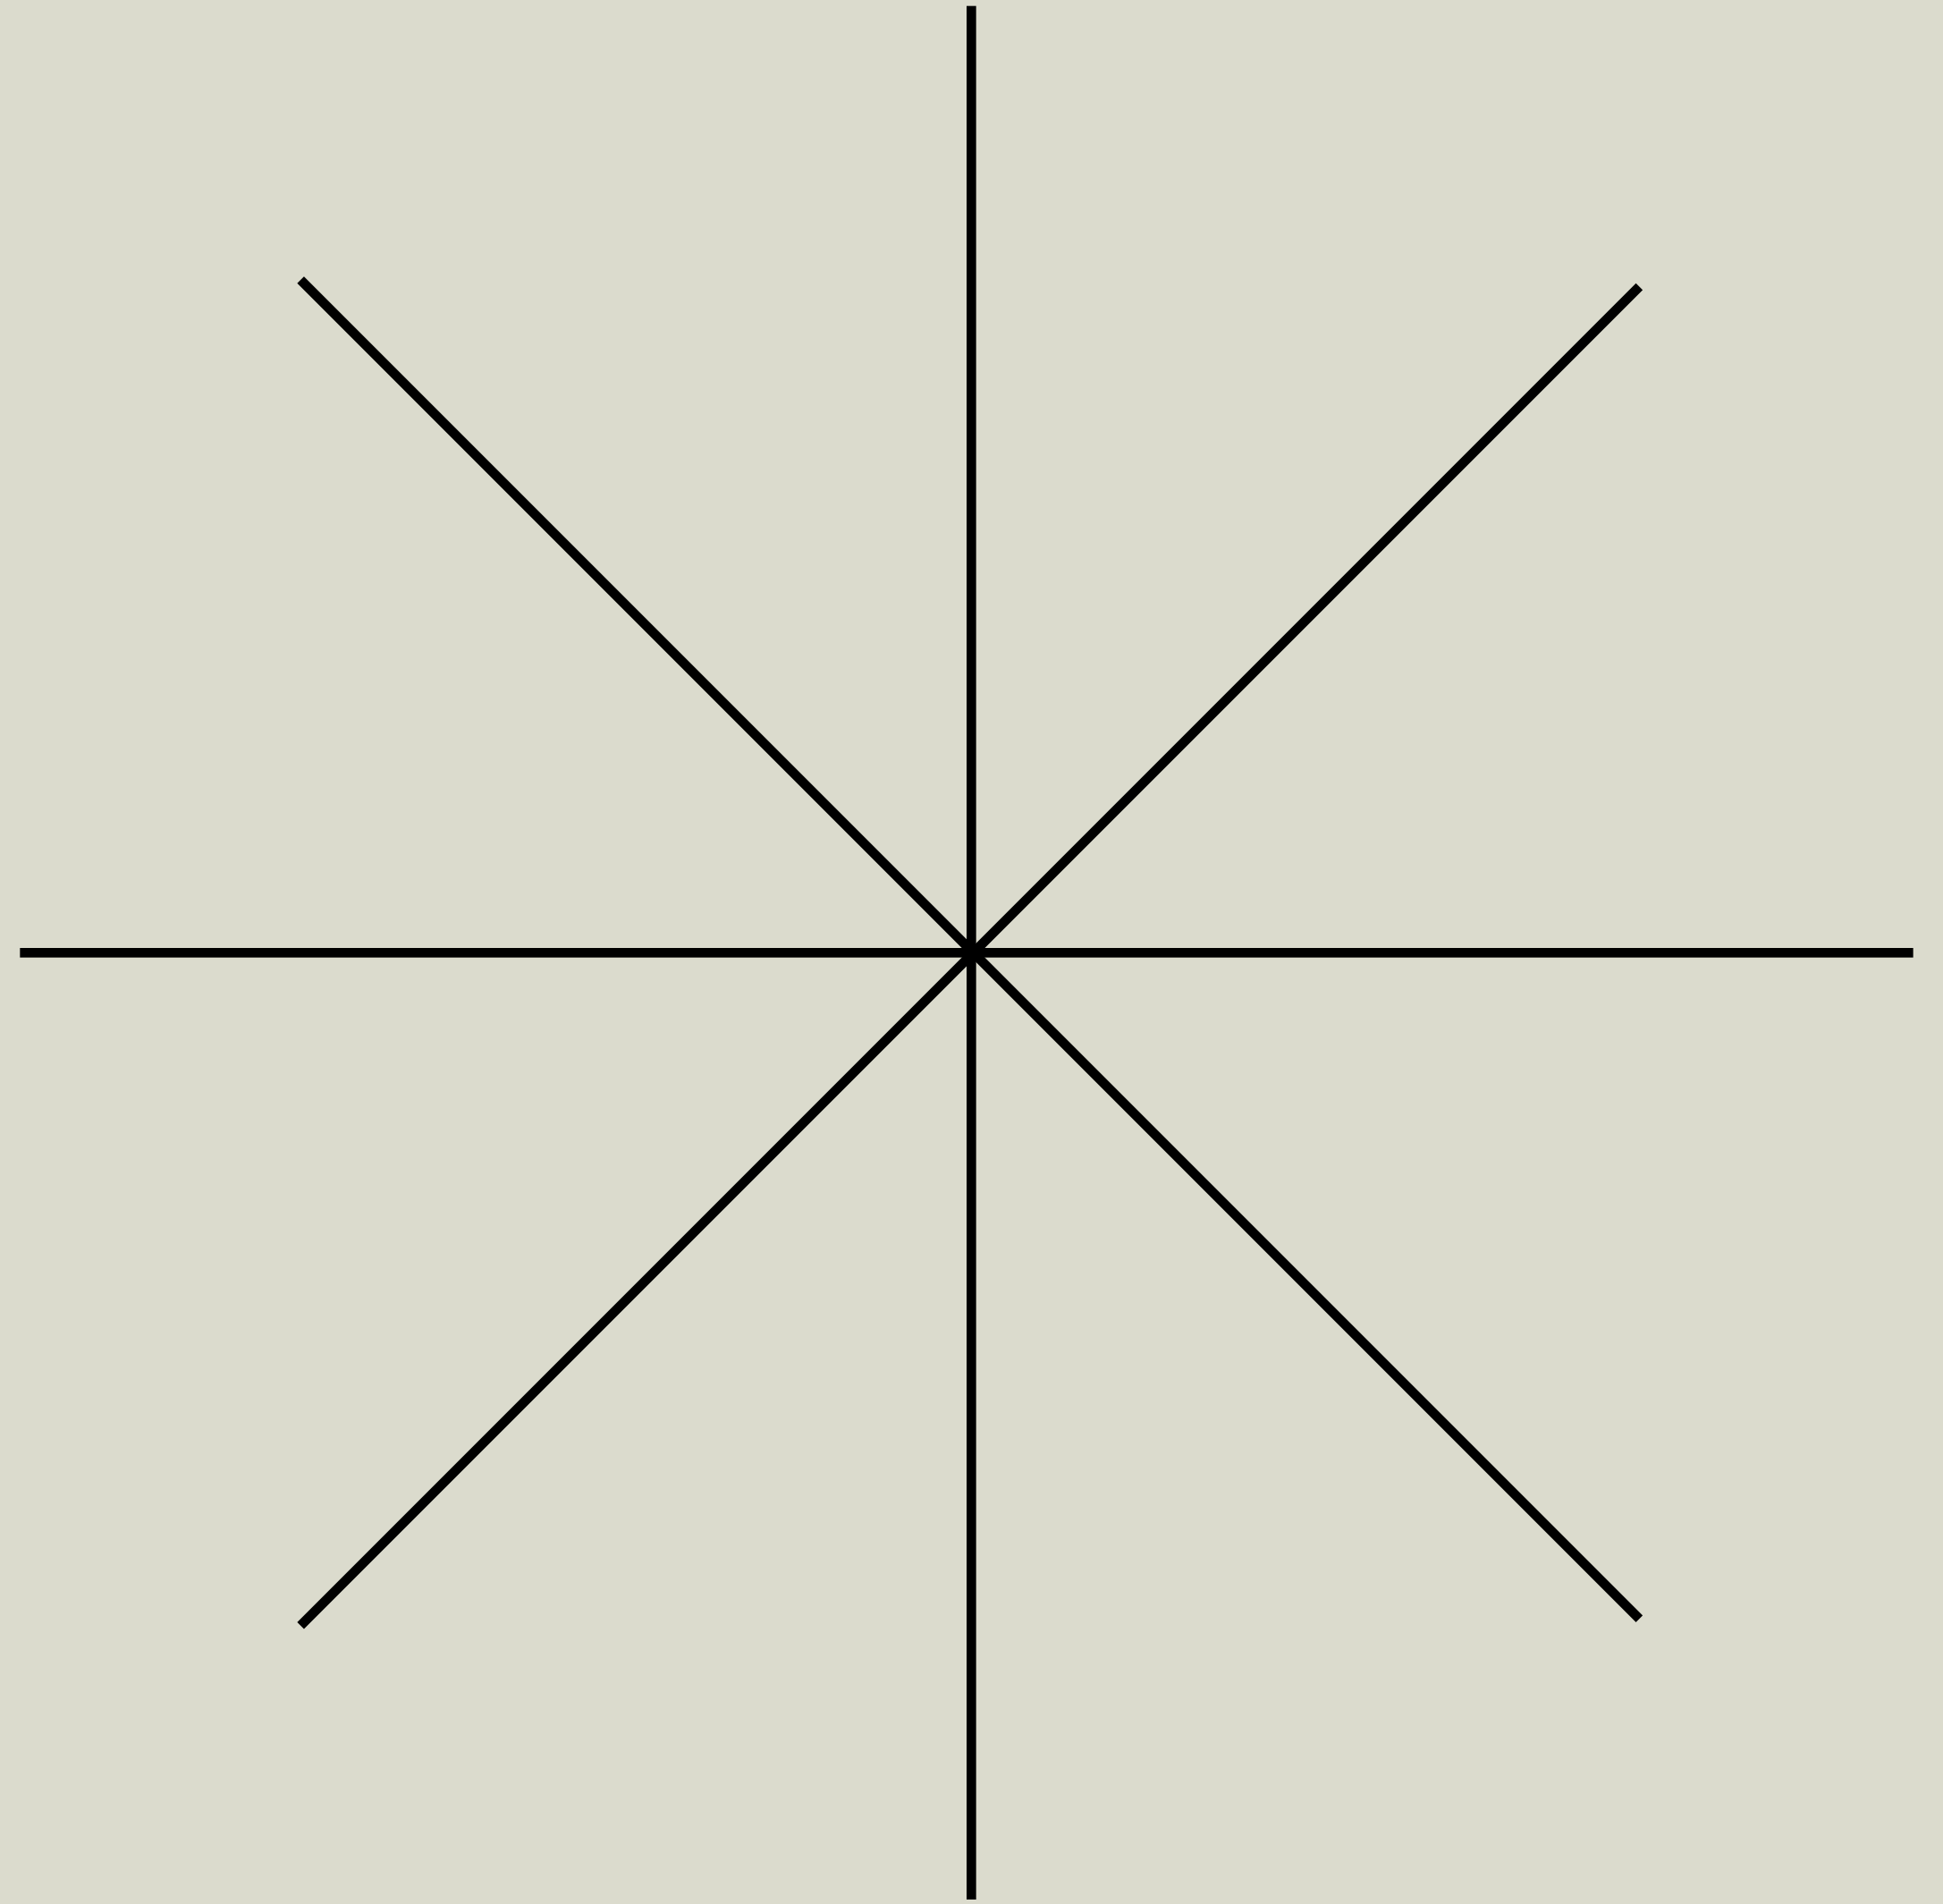
✓ Strong semantic compliance

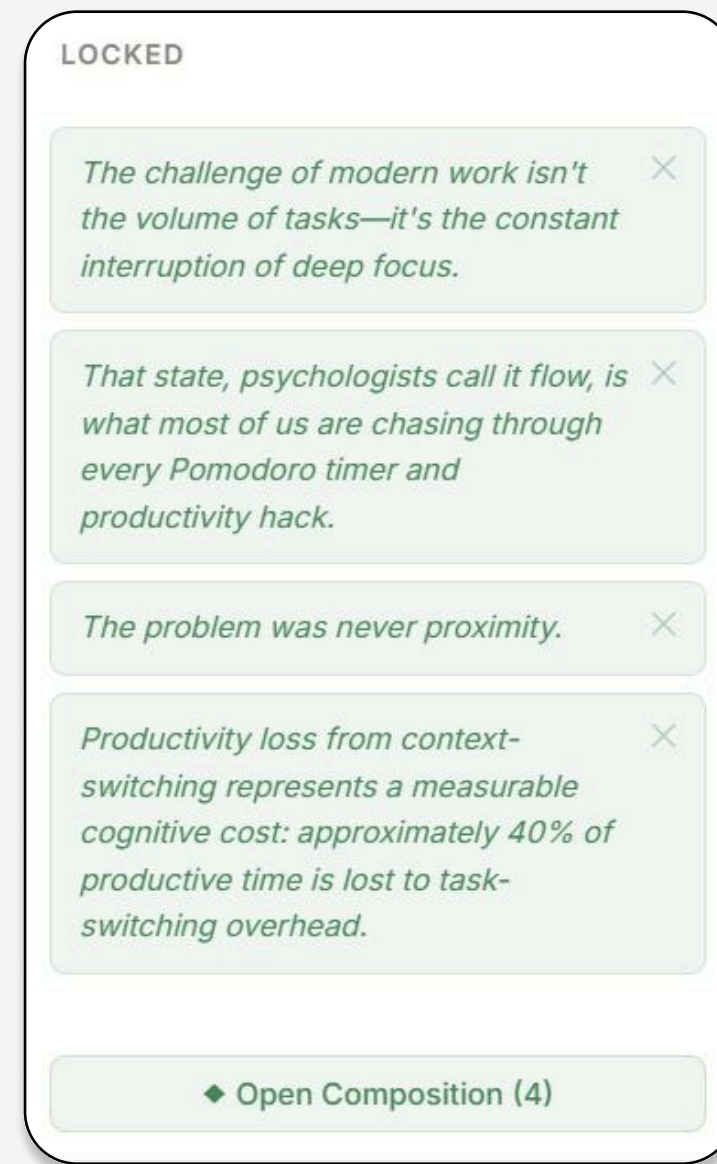
The model understands and respects the intent of the lock.

⚠ Weak syntactic compliance

Punctuation or phrasing may be subtly altered.

05 REFLECTION





Evaluating non-deterministic systems

Known Limitations

SAME-MODEL BIAS

Claude both generated outputs and judged coherence. The judge may favor its own writing style.

SINGLE PROMPT TESTED

These results may not be consistent when varying other prompts that are given, or if different types of sentences are locked.

SMALL SAMPLE

15 runs is pretty small! I kept this small for the sake of feasibility, but would definitely want to complete a much more exhaustive evaluation process if I were to take this project further.

SCOPE CREEP IS REAL

This project was a lesson in scope management. What started as a simple canvas for generating text variations gradually expanded into a system with style tags, dimension sliders, sentence locking, and a composition view. The evaluation process followed the same pattern. At one point the design included 3 prompts × 20 runs across multiple locked element positions before being scoped down to something focused and meaningful. Working with generative AI amplifies this tendency: because the model is capable of so much, the temptation to keep extending the system is constant. I found that the deeper design challenge isn't learning to prompt the model effectively, it's developing the discipline to constrain the scope of what you ask it to do. I was grateful to have the chance to tackle something like this, and hope to keep expanding on this project in the future.

Thank You!