




Designing for Control & Steering

Beyond Chat: designing
control into agentic AI

Agenda

■ 01 Diagnosis

Why chat breaks

■ 02 Framework

Control-aware UX


■ 03 Patterns

Beyond Chat, live


THE MOMENT

AI feels like the early web again

Powerful technology. Few stable conventions.



SAME MOMENT,
NEW MEDIUM
→



1998 — the web. Powerful technology, no conventions yet.


2026 — AI. Raw capability, conventions still unwritten.

The conventions aren't written yet — that's the opening to shape them.

THE MOMENT

Designers vibe code. Developers vibe design.

The boundary used to be the IDE. That boundary is going.



Same work. Increasingly the same vocabulary.
What stays unique is judgment — and judgment needs control.

01

Diagnosis







Why chat breaks



DIAGNOSIS

Six tools. One pattern.


A blinking cursor in a box.



DIAGNOSIS

Chat is a command bottleneck

Chat is powerful. But it compresses too much.




Everything is translated into sequential text.

DIAGNOSIS

AI becomes a Black Box

Lack of Transparency




Loss of Judgment

Loss of Control

DIAGNOSIS

The slot-machine moment

When AI fails, users often pull the lever again.



Faster output can hide weaker understanding.

Good AI UX should not train users into passive retries.

DIAGNOSIS

The cost is measurable.

If we delegate thinking to AI, our own thinking measurably weakens.

47%

less brain activity in AI-assisted users
compared to peers writing unaided.

83%

couldn't later recall key claims
in their own AI-written texts.

"The user becomes a spectator of their own output."

“The AI we want to see must be
human-centered and dynamic.”


— Dan Saffer

02 Framework


Control-aware UX




The Solution: A Systematic Three-Layer Model



Keeping people in command



Generative AI
Creates the space of options — fast, and many.
But options aren't decisions.



YOU
Decide what ships — authorship & judgment stay human.

Control is designed in, not bolted on — a first-class design goal, equal to capability.



Gradual Autonomy



Human Control

Human acts manually
and performs the
detailed work

Human-AI Collaboration




AI Automation

AI acts autonomously
and takes over mass
production

Automation is a design decision

It is not chosen by capability. It is chosen by risk.



The question is not: What can the AI do?

The question is: **What happens if it is wrong?**

QUICK QUIZ

You add autonomy to an AI feature. What should decide how much?

A How capable the model is


B How reversible the action is

C How much time it saves

D How good it looks in the demo

Vote in the app


Control is not the opposite of automation



Control means the user can:

- See what the AI is doing
- Understand when to trust it
- Intervene before execution
- Recover when it fails
- Stay capable over time


Implementation of Control-Aware UX



Chat

Sequential chat & command pattern.
To change one thing, you describe everything again.


Beyond-Chat
& XUI



Controlled Implementation

Visual, directly manipulable parameters.
Change one thing — only that thing changes.

DOPE: Four dimensions of Control-Aware UX



03 Patterns


Beyond Chat, live



PATTERNS

60 live patterns — beyondchat.design

A public catalog of working prototypes — not slides.



The Landscape: Five Lanes of Beyond-Chat Patterns

From a public catalog of interactive patterns I curate.

INLINE & CONTEXTUAL15



Replace brittle prompt wording with direct, local controls inside the flow of work.

VISUAL & GENERATIVE9



Turn outputs into visual, manipulable surfaces instead of static prose responses.

AGENT WORKFLOW17



Plan, intervene, and supervise multi-step agent runs from start to finish.

TRUST & VERIFICATION9



Surface uncertainty, sources, and reasoning so users can verify before they trust.

MEMORY, CONTEXT & BOUNDARY10



Make what the AI remembers, sees, and is constrained by visible and editable.

Five lanes, one framework — DOPE. [Let's look at one from each, live.](#)

● LIVE DEMO

DEMO 1/5

Property Panel

From prompt to parameters.

LANE **01 · Inline & Contextual**

DOPE FOCUS

D O P E

live at beyondchat.design

PROPERTY PANEL · PATTERN 01

16:9



ASPECT RATIO

1:1

4:5

16:9

LOOK

RAW

CINEMATIC

ILLUSTRATION

ON-IMAGE TEXT

STYLIZE

250

● LIVE DEMO

DEMO 2/5


Spatial Canvas

Thinking in space, not a scroll.

LANE 02 · Visual & Generative

DOPE FOCUS D O P E

live at beyondchat.design



● LIVE DEMO

DEMO 3 / 5

Editable Plan

Agents need supervision before execution.

LANE 03 · Agent Workflow

DOPE FOCUS D O P E

live at beyondchat.design

■ ■ ■ EDITABLE PLAN

1

Audit current component library

EDIT

2

Draft token migration map

EDIT

3

Apply tokens across 14 screens

COMMENT

+ Add a step

Approve & run

● LIVE DEMO

DEMO 4/5

Citation Trail

Trust needs evidence,
not confidence.

LANE 04 · Trust & Verification

DOPE FOCUS D O P E

live at beyondchat.design

CITATION TRAIL

Adoption rose 37% in the first quarter, driven mostly by teams already using design tokens . Retention held steady across both cohorts.

Source A

Source B

Source C

● LIVE DEMO

DEMO 5 / 5

AI Memory

What it knows, you
can see and edit.

LANE 05 · Memory, Context & Boundary

DOPE FOCUS D O P E

live at beyondchat.design

AI MEMORY

ALL

PREFERENCES

PROJECTS

Prefers concise, structured answers

PREFERENCE

Works in design systems & UX research

CONTEXT

Prefers Figma & geometric type

PREFERENCE

Old role: front-end developer

OUTDATED

PATTERNS

From Catalog to Practice

Five lanes. One framework — DOPE.

INLINE & CONTEXTUAL15



Property Panel

DIRECTABILITY

VISUAL & GENERATIVE9



Spatial Canvas

DIRECTABILITY & OBSERVABILITY

AGENT WORKFLOW17



Editable Plan

DIRECTABILITY & PREDICTABILITY

TRUST & VERIFICATION9



Citation Trail

DEMO 3 · EXPLAINABILITY

MEMORY, CONTEXT & BOUNDARY10



AI Memory

PREDICTABILITY

PATTERNS

More from the catalog

Beyond the five live demos — more in every lane.


INLINE & CONTEXTUAL15  Ghost Text Tone of Voice Region Lock	VISUAL & GENERATIVE9  Generative UI Parallel Variants Real-time Gen	AGENT WORKFLOW17  Agentic Dashboard Autonomy Dial Approval Gates	TRUST & VERIFICATION9  Reasoning Trace Confidence Signals AI Provenance	MEMORY, CONTEXT & BOUNDARY10  Source Toggle Persona Switcher Constraint Rules
---	--	--	---	---

Explore the full catalog: beyondchat.design

PATTERNS


You already use these patterns

Beyond-chat patterns ship in tools you open every day




Perplexity

Citation Trail — see the sources



Cursor

Editable Plan — edit before it runs



ChatGPT

AI Memory — see what it knows

What separates these tools isn't the model — **it's the interface.**

TAKEAWAY

From Prompting to Control & Steering

What Beyond-Chat is really about

THE CORE

Beyond Chat

Canvas, panels,
dashboards — not
just a text box.

Steerability

Direct and intervene
in real time — not re-
prompt.

Trust

See, verify, and predict
what the AI does.

Sovereignty

People stay in
command — subject,
not object.

Chat is where AI started. Not where it ends.

The interface beyond the chat box is still unwritten —
and it's ours to design.

beyondchat.design

Designing for Control & Steering

Q&A



Dominik Witzke

[linkedin.com/in/dominikwitzke](https://www.linkedin.com/in/dominikwitzke)

dwitzke@star.global