

Open in app ↗

 Search Write

All your favorite parts of Medium are now in one sidebar for easy access.

Okay, got it

Design Principles for AI



Dan Saffer

Follow

7 min read · 5 days ago



As we've been working through our AI use case examples and research, it made sense to create some Design Principles to guide our work, both for AI in general and the individual use cases in particular.

Design principles are more than guidelines. They define success and set the direction for our work. They act as north stars, guiding us through ambiguity and helping us make better choices.

The process of writing these principles was exciting because we could finally give shape to burgeoning opinions around AI. However, it can be frustrating if a team ends up following rules they don't believe in. That's why we went through multiple rounds of iteration and group revision until the principles felt truly ours.

As we explored AI use cases and conducted research, we began to shape our own vision of how AI should evolve. Once we understood what already

exists, we returned to the question that sparked this project in the first place:

Why are current AI design paradigms missing the mark? Our answer was

All your favorite parts of Medium are now in one sidebar for easy access.

nt to see in the world must be human-centered and does that really mean?

From getting lost in endless sketching and prototyping, we developed a set of Design Principles. We aim to use these principles to ground our work moving forward and keep us aligned with our greater purpose of ensuring UI for AI serves human needs and values.

GENERAL DESIGN PRINCIPLES

The UI for AI Team

Control and Agency

Prioritize user intent over AI output. The AI should adapt to user intent, supporting their decisions rather than making them. For example, a designer can edit a header or merge layouts while keeping the rest intact. (Zeana)

Emphasize user agency. Enable users to constrain system defaults. Consider a user choosing which notes the AI retains or which interactions are logged without disruptions to their workflow. (Erica)

Handling AI Uncertainty

It's more than just an algorithm. Account for how users interpret, trust, and act on outputs. (Claire)

Make uncertainty visible and actionable. Guide users to interpret and recover from errors even when the AI can't tell it's wrong. (Adithi)

Understanding AI Capabilities and Limitations

Deliver on needs, not necessarily commands. A user's command might be incomplete, or even counterproductive. But underneath lies a desire to understand. AI should interpret intent, context, and desired outcomes, then deliver on the user's underlying purpose. (Valerie)

All your favorite parts of Medium are now in one sidebar for easy access.

Design interfaces and flows that make the system's strengths and limitations obvious, so users can set realistic expectations. (Dave)

Make the next step clear. Spawn tools based on behavior and intent. (Dan)

Contextual History

Design for continuity. Ensure users can easily view, resume, and carry context across tasks. For example, resurface sources a user previously explored to connect past insights with new findings. (Hanara)

Flexible Workflows

Support adaptability over rigidity. Interfaces should flex to users' changing goals and contexts, enabling multiple pathways instead of enforcing a single way to work. For example, switching views for different interactions: a linear chat view for conversation, a map/outline view for idea exploration. (Celine)

Iterative Exploration

Design for exploration and reflection. Support both the exploration of new alternatives and the ability to build on past ideas. For instance, allow users to annotate or bookmark key moments as they experiment. (Holly)

Design for reassurance. Let AI guide with clarity and empathy so users feel supported, not second-guessed. (Yukti)

All your favorite parts of Medium are now in one sidebar for easy access.


Clarity. Make it easy for the AI to foster creativity without

Trust and Clarity

Provide clear communication of the actions the AI performed. (Cole)

Reveal the process. Demonstrate how the AI thinks and works to help users understand how to become better at using it. (Jason)

USE CASE SPECIFIC DESIGN PRINCIPLES



AI Use Cases diagram by Celine Tseng



All your favorite parts of Medium are now in one sidebar for easy access.

iki

System's Scope. Communicate what the AI can and cannot successfully achieve. Example: At the start of a session, the system makes its scope clear: "I can do X or suggest Y, but not Z." This sets boundaries while still encouraging creative exploration.

Automate tasks while preserving human control. Provide scaffolding and automation, while preserving the user's creative control and allowing for complete freedom of expression. Example: The system offers a few starting layouts based on the user's prompt, and the user can move, swap, or delete any element. This accelerates brainstorming while the person stays in charge of shaping the final result.

Frame early actions as immediate momentum. Designing the initial interaction to feel like progress, not a start from a blank canvas. Example: The system provides starting points such as templates, sample structures, or draft variations, so users can react and build right away

Conversation Flow and Prompt Editing

Celine and Holly

Enable flexible and diverse options for refining prompts. Prompt editing should integrate naturally into conversation, should support diverse methods of prompt iteration, including both user-driven and AI-assisted approaches. Example: Users can iterate prompts through text editing, UI buttons, AI-assisted suggestions, or a combination of these methods.

Emphasize divergent exploration over forced linearity. Conversations

should accommodate branching thought processes and iterative exploration at structures. Example: A tree-like view lets users fork off to explore different paths, and move fluidly between ideas for iteration.

All your favorite parts of Medium are now in one sidebar for easy access.

Enable precise access along with high-level navigation. Users should be able to scan the conversation flow at a high level and jump directly into specific points for precise edits. Example: A zoomable map shows the overall conversation at a glance, while clicking any node opens its prompt for more granular edits.

Refinement Flow

Celine, Zeana, Erica

Support multi-level refinement from overview to fine details. Users should be able to adjust designs at different levels, from big-picture layout changes down to small tweaks like colors or text. Example: Users can refine at different levels — rewriting a whole paragraph for tone, adjusting a single sentence or word for clarity, or combining layout from one option with style from another.

Enable precision while preserving structure. Edits should target exactly what the user wants to change, while keeping the overall hierarchy and layout intact. Example: Users highlight multiple words simultaneously in a generated output (ex: “run,” “eat,” and “said”), the system revises synonyms while preserving structure and swapping multiple elements in the next output.

Emphasize co-authorship over replacement. AI should act as a collaborative partner, refining and building on user ideas instead of discarding them and

All your favorite parts of Medium are now in one sidebar for easy access.

Example: Inline suggestions instead of overwrite — Show AI or side-by-side alternatives, so the user can accept, reject, or than losing their original text.

Selection and Prioritization

Claire and Valerie

Adapt to changing goals by filtering what matters from complex contexts. Enable AI to adapt priorities in real time by distilling complexity into clarity, building user trust in its decisions.

Design for a person, not just a persona. A person isn't static. Understanding the person means being attuned to situational signals rather than assuming their priorities based on preconceived notions.

Empower precise expression of needs to ensure alignment. AI should give people the tools and scaffolding to articulate their needs, guiding them even if they don't know the "right" technical terms.

Multitasking and Context Switching

Hanara and Holly

Preserve context so users can easily resume tasks. Systems should remember and represent the task flow with clear markers, so users don't have to reconstruct their working memory each time they return. Example: When users return to a draft, the system reopens at their last edit with changes visible, so they can immediately pick up where they left off.

Maintain clear boundaries between different tasks. Interfaces should support compartmentalized task spaces, ensuring that progress and context doesn't interfere with another. Example: If someone is working on Project A and another project in parallel, switching to Project B should show only Project B's progress while keeping reminders and comments from Project A.

All your favorite parts of Medium are now in one sidebar for easy access.

Anchor workflows in chronology to track fragmented flows.

Use temporal structures and visible anchors to help users trace, reenter, and orient themselves within their workflows. Example: A design tool lays out a simple timeline — previous mockups, newly integrated feedback — so users can quickly see the flow of work belonging to that project.

Error Detection, Reporting, and Awareness

Yukti and Adithi

Prevent user errors before they occur. Anticipate and flag potential input issues in real time so users can correct mistakes early and stay in flow. Example: When a prompt lacks context, flag it before submission instead of returning a vague, unhelpful response.

Surface system uncertainty before it misleads. Be transparent when the AI may be unsure or incomplete in its responses, helping users understand potential limitations and make informed decisions. Example: When an AI generates a confident-sounding answer with weak evidence, add — “I’m filling some gaps here, double-check important details.”

Guide iterations with actionable refinements. AI regeneration can produce unpredictable results. Offer targeted adjustment options to help users

improve outputs rather than relying on “Try again” options. Example: When an AI writes a formal email but a casual tone is intended, surface options like “casual” instead of only “Retry.”

All your favorite parts of Medium are now in one sidebar for easy access.

ig

Dave, Erica, Jason

Connect conversations, don't archive them. Memory should connect threads across conversations, not log everything that happens. Focus on what helps users maintain flow and context, not comprehensive data storage. Example: When returning to a project discussion, surface key decisions and open questions that enable continuation, rather than chronologically listing every project mention.

Make memory visible to make it trustworthy. Users must see what's remembered, when it influences responses, and why it was important. Invisible memory feels manipulative; visible memory feels collaborative. Example: Show “Based on your preference for concise updates...” with an option to view or edit that memory, rather than silently adjusting response style.

Design forgetting as carefully as remembering. AI should gracefully sunset, terminate, or reduce irrelevant or outdated information. Complement it with condensed, efficient storage of simple facts prioritized. Example: Project-specific details automatically become less prominent after completion, while communication preferences persist unless explicitly changed.

Preserve and augment human meaning, don't dilute it. Memory should strengthen users' context-making and recall while preserving their sense of

authorship. Enhance human meaning rather than summarizing it away.

Example: Maintain “I’m nervous about this presentation but excited about
in reducing it to “presentation anxiety = moderate.”

All your favorite parts of
Medium are now in one
sidebar for easy access.

ie introduction to these principles.

AI

UI

Design Principles

UX

Hci



Published in UI for AI

Follow

353 followers · Last published 5 days ago

Explorations around creating new pieces of UI optimized for AI-powered products



Written by Dan Saffer

Follow

16.7K followers · 262 following

Designer. Product Leader. Author. Professor.

Responses (2)



Harry Baya

What are your thoughts?

All your favorite parts of Medium are now in one sidebar for easy access.



- The design principles are not very design-centric. You could easily have 80% of these to be "product guidelines for AI". Think about scoping them more towards design. Yes, product is also design, but design contributes uniquely... [more](#)



16

[Reply](#)**Ana Luísa Ponsirenas**

5 days ago




These are so good. I specifically like the refinement flow and the forgetting principles.



11


[Reply](#)

More from Dan Saffer and UI for AI



In UI for AI by Dan Saffer

Welcome to UI for AI



In UI for AI by Dan Saffer

Diving Deep into AI Use Cases


AI needs UI. AI makes great UI more important than ever so that we can...

This week, the UI for AI team explored their use cases, seeking design opportunities and...

All your favorite parts of Medium are now in one sidebar for easy access.

Sep 21  52  



 In UI for AI by Yukti Poddar

Designing the Un-Chatbot: Key Insights from Our Research

After our initial brainstorming session, our study group dove headfirst into the frontier ...

Sep 15  64  

Jul 21  413  18  



 Dan Saffer


AI Needs UI

It seems like every day, someone who doesn't know anything about design proclaims "UI is..."

See all from Dan Saffer

See all from UI for AI

Recommended from Medium



05 Help When People Get Stuck

Surface agent suggestions during error states, not randomly. When Warp users hit Git errors, they auto-suggest "Let agent fix this" with one click. Catch frustration moments, don't interrupt flow.

06 Start Small, Grow Trust

Warp began with simple terminal commands, then added file editing, web reading, and system access. Each month brought new capabilities as users got comfortable with agent reliability.

07 Don't Let Power Users Bankrupt You

Zach admits Warp loses money on heavy users under fixed subscriptions. A power user costs \$80

08 Price Like Cell Phone Plans


Use base subscription plus overages. Give predictable costs for normal usage but protect

 In UX Collective by Simon Robinson 

Start with the experience

What Genie taught us about customer-centric innovation—and how those lessons still...

Sep 26  171  2  



 In Bootcamp by Mohan Kumar S

UX Trends to Watch in 2026: Designing for Empathy,...

As we count down the years to 2026, UX design is no longer just “good UI” or “usable...


★ 6d ago  66  2  

 Aakash Gupta

The \$1M ARR Every 10 Days Playbook: How Warp Cracked the...

Most companies are still figuring out how to make AI profitable. Warp isn't one of them.

5d ago  42  




 Hriday Checker


Why everything looks the same in 2025

The Copy-Paste World

★ Sep 23  730  30  



All your favorite parts of Medium are now in one sidebar for easy access.



 In UX Planet by Amber Krishan

AI-Native? Future-ready?—Go Design-Led first

To unlock real impact from AI, data, or emerging tech—design must lead.

Sep 24  1



 In UI for AI by Dan Saffer

Diving Deep into AI Use Cases

This week, the UI for AI team explored their use cases, seeking design opportunities and...

Sep 21  52



See more recommendations