

DesignForge

Design Critique & Research Studio

Onboarding Manual

Version 1.0 — April 2026

Screen Edition — Dark Theme

C O N T E N T S

01	Before You Begin <i>The philosophy beneath the interface</i>	5
02	First Light <i>Installation, backends, and your first connection</i>	7
03	The Source Vault <i>Uploading, organizing, and weighting your materials</i>	10
04	Lenses <i>Choosing how to see your own work</i>	13
05	The Critique Engine <i>Asking the right questions, reading the answers</i>	16
06	Deep Dives <i>When the surface isn't enough</i>	19
07	Thoughts & Notes <i>Your private sketchpad inside the studio</i>	21
08	History & Export <i>Keeping what matters, letting go of the rest</i>	23
09	Settings & Customization <i>Making it yours</i>	25
10	The Interface in Detail <i>Every panel, every button, every quiet corner</i>	27
11	Keyboard Shortcuts & Tips <i>Working faster without thinking harder</i>	30
12	Troubleshooting <i>When things go sideways</i>	32
13	Design Philosophy <i>Why DesignForge exists at all</i>	34
14	Appendix: Color & Typography Tokens <i>The aesthetic DNA</i>	36
15	Real-World Workflows <i>Five scenarios from the trenches</i>	39
16	Advanced Techniques <i>Power moves for the committed</i>	42
17	Collaboration & Sharing <i>Working together, apart</i>	44
18	Ethics of AI Critique <i>What the machine can't do</i>	46
	Quick Reference <i>Everything at a glance</i>	48
	Glossary <i>Terms worth knowing</i>	49

Before You Begin

You already know more about design than you think. DesignForge isn't here to teach you taste—it's here to give your taste a sharper tongue.

Picture the last time you looked at a design—a website, a poster, a product label—and felt that small tug in your chest. Something was off. Or something was quietly right. You couldn't always name it, but you felt it. DesignForge lives in that gap between feeling and naming.

It's a critique studio, not an answer machine. You bring the work, you bring the questions, and the tool helps you interrogate both through a set of configurable lenses—each one a different way of seeing. Think of it as a panel of thoughtful colleagues who never get tired, never get defensive, and never check their phones mid-conversation.

But let's be honest with each other for a moment. No tool replaces the slow, unglamorous work of looking. Of sitting with discomfort. Of admitting that the layout you spent three days on might need to be torn down and rebuilt. DesignForge can accelerate insight, but it can't manufacture it. The insight has to come from you, from the friction between what you intended and what actually landed.

What This Manual Is (And Isn't)

This is a walking companion, not a reference tome. We'll move through the interface panel by panel, but we'll also linger where it matters—in the philosophy behind the design, in the small choices that compound into something meaningful. If you want the quick-start version, jump to Chapter Two. Nobody will judge you.

If you're the type who reads the liner notes before playing the album, stay here. We have things to talk about.

Who This Is For

Designers who critique their own work. Researchers who need to synthesize findings through multiple frameworks. Students learning to see. Creative directors who want a structured second opinion at 2 AM when the team has gone home. Anyone who has ever stared at a screen and whispered, “Is this good? Is this actually good?”

You don't need to be technical. You don't need to know what an API key is—though you'll learn, gently, in the next chapter. You need curiosity and a willingness to let a tool ask you uncomfortable questions about your own work.

First Light

Every great tool has a moment where it wakes up. For DesignForge, that moment is a small green dot.

Opening DesignForge for the first time feels like walking into a dim, well-organized studio. The dark interface isn't an aesthetic choice for its own sake—it's designed to keep your eyes on the work, not the chrome. The muted purples and soft borders recede, leaving your content in the foreground. If you prefer daylight, there's a light mode. We'll get to that.

Choosing Your Backend

DesignForge doesn't think on its own. It needs a language model—a brain, if you want the anthropomorphic version—running somewhere. You have four choices, and each one trades off something different:

Ollama
Free, private, runs on your own machine. Download from ollama.com, pull a model, and you're live. Your data never leaves your laptop. The tradeoff is speed and capability—local models are smaller, sometimes rougher.

Anthropic
Claude's API. Sharp, nuanced, excellent at critique. Requires an API key and costs money per use. Your data travels to Anthropic's servers. The tradeoff is privacy for quality.

WebAI
Browser-based local inference. Experimental, fascinating, sometimes unreliable. For the curious and the patient.

StudioLM
An Ollama-compatible server you run on a different port or machine. For studios with a dedicated inference box.

To set up Ollama—the default—open your terminal and type two commands. The first installs the server. The second pulls the model that does the thinking. Once both are running, switch to DesignForge. That small dot in the top-right corner? It should be green now. If it's yellow, it's still looking. If it's red, something went wrong—but we'll troubleshoot that in Chapter Twelve.

The Green Dot Moment

There's a small, specific satisfaction in seeing that dot turn green. It means the circuit is closed—your studio is connected to its brain, and the critique engine is ready. Take a breath. The real work starts now.

If you chose Anthropic, head to the Settings tab and paste your API key. The model field will default to Claude Sonnet. If you've got access to other Claude models, type the model string in directly. DesignForge doesn't gatekeep—it'll try whatever you tell it to.

Your First Look Around

The interface is a single screen divided into a sidebar and a main panel. The sidebar holds your Thoughts—we'll get to those—and a search bar. The main panel has seven tabs, each a different room in the studio:

Sources — Where your materials live. Briefs, research, brand guidelines.

Lenses — The frameworks you'll use to examine your work.

Critique — Where the engine reads your sources through your lenses.

Deep Dive — For when you need to go further into one dimension.

Notes — Your private scratchpad. Timestamped, searchable.

History — Every critique and deep dive, stored and exportable.

Settings — Backend configuration, theme toggle, data management.

You'll move between these tabs constantly. They're not a sequence—they're rooms in a studio, and you'll develop your own circuit through them.

The Source Vault

Your sources are the raw material of critique. The quality of what you put in determines the depth of what comes out.

The Source Vault is the first tab for a reason. Before DesignForge can say anything meaningful about your work, it needs to understand what your work is, what it's responding to, and what constraints it's operating within.

You upload plain text or Markdown files—design briefs, user research transcripts, brand guidelines, competitor analyses, meeting notes, whatever forms the constellation around your design decisions. Drag them into the upload zone, or click to browse.

Three Folders, Three Trust Levels

Every uploaded file lands in one of three folders. Think of them as levels of authority:

Primary Sources

The core documents. Your design brief, the user research, the brand bible. These carry the most weight when the critique engine reasons about your work. If your brief says the target audience is elderly users, the engine will hold that truth close.

Reference Materials

Supporting documents. Competitor analyses, style guides from adjacent brands, articles you found useful. The engine reads them, but treats them as context rather than commandment.

Scratch Notes

Your own observations, hunches, loose threads. Lower weight, but sometimes a note you scribbled at midnight turns out to be the sharpest thing in the vault.

Each folder has a weight slider. By default, Primary Sources are weighted highest. But you can adjust this—maybe you're in a phase where the competitor landscape matters more than the brief. The sliders let you shift the center of gravity without removing anything.

Indexing: Teaching the Engine to Read

Once your files are uploaded, hit the "Index All Sources" button. This creates embeddings—mathematical representations of your content that let the engine find relevant passages when you ask a question. It's like building an index for a book, except the book is everything you've ever uploaded.

Indexing uses the embedding model specified in Settings (by default, nomic-embed-text for Ollama). It runs locally unless you've chosen a cloud backend. The process is quiet—you'll see a brief status update, then it's done.

A Note on What to Upload

Resist the urge to upload everything. The critique engine works best when the source vault is curated, not stuffed. Three precise documents will yield better insight than thirty vague ones. Think of it this way: if you wouldn't bring a document into a real critique session, don't bring it here.

Clarity in, clarity out. Noise in, noise out. The vault rewards intentionality.

You can always clear sources and start fresh. The engine doesn't hold grudges. It doesn't remember previous sessions unless you explicitly save state. Every session can be a clean slate if you want it to be.

Lenses

A lens doesn't change what's there. It changes what you see.

This is the heart of DesignForge. Lenses are the frameworks through which the critique engine examines your work. Each lens asks a different question, privileges a different value, notices different things.

Imagine you've designed a landing page. Through a typography lens, the engine might notice your line height is too tight for body copy on mobile. Through an accessibility lens, it might flag insufficient contrast. Through a narrative lens, it might question whether the visual hierarchy supports the story you're trying to tell. Same page, different eyes.

Preset Packs

DesignForge ships with several preset packs—curated collections of lenses designed for specific critique styles:

Critic

The default. A balanced set of lenses covering visual hierarchy, typography, color theory, spacing, and consistency. Good for general-purpose critique.

Researcher

Optimized for user research synthesis. Lenses focused on methodology, bias detection, pattern recognition, and insight extraction.

Brand Auditor

For evaluating brand consistency. Lenses examining voice, visual identity, emotional resonance, and strategic alignment.

Systems Thinker

For design systems work. Lenses focused on component reuse, naming conventions, scalability, and documentation quality.

Building Your Own Lenses

The preset packs are starting points, not destinations. The real power is in crafting your own lenses—questions that are specific to your work, your client, your obsessions.

A lens has two components: a name and a core question. The name is a label—"Mobile Readability" or "Emotional Pacing" or "Does This Spark Joy." The core question is what the engine asks of your work when that lens is active. Make it specific. Make it pointed. "Is the typography good?" is a weak lens. "Does the type scale create a clear reading rhythm across breakpoints?" is a strong one.

Intensity and Toggling

Each lens has an intensity slider and a toggle. The slider controls how much weight the engine gives that lens in the critique—think of it as volume control. Crank up typography, dial down color, and the critique will spend more words on letterforms. The toggle turns a lens on or off entirely without deleting it. You're building a palette of perspectives, and not every perspective needs to speak every time.

The best critiques come from tension between lenses. Don't make them all agree.

Here's a secret: the most interesting insights happen when your lenses disagree. An accessibility lens might demand larger text. A visual density lens might want tighter spacing. The friction between those demands is where design decisions live. Don't resolve the tension—let the engine surface it.

The Critique Engine

Ask a vague question, get a vague answer. Ask a question that makes you slightly uncomfortable, and you might learn something.

The Critique tab is where everything converges—your sources, your lenses, your questions. You type a prompt, the engine reads your source vault through every active lens, and writes a structured critique.

The quality of the critique depends almost entirely on the quality of your prompt. This isn't a search bar. It's a conversation opener. The best prompts are specific, contextual, and slightly brave:

"Does our visual hierarchy support the core user task, or are we decorating around it?"

"What assumptions about our users does this layout encode? Are they fair?"

"If this brand system were a person, what would their voice sound like at small sizes?"

"Where does the design ask the user to do cognitive work that the interface should be doing?"

Notice what these prompts have in common. They're not asking "Is it good?" They're asking the design to defend itself. They're creating a context where the engine has to take a position, not just describe.

Reading the Output

The critique arrives as structured prose—not a checklist, not a score. It reads like a letter from a thoughtful colleague who has spent time with your materials. The engine will reference specific sources from your vault. It will address each active lens. It will sometimes contradict itself, because good critique acknowledges complexity.

Don't read it looking for validation. Read it looking for the sentence that makes you pause. That's where the insight is.

The Critique Workflow

A typical critique session looks like this:

1. **Upload and index**
Get your sources into the vault. Index them.
2. **Choose your lenses**
Activate the perspectives that matter for this critique.
3. **Write your prompt**
Be specific. Be brave.
4. **Run the critique**
Hit the button. Wait. The engine is reading and thinking.
5. **Read slowly**
Don't skim. Read the way you'd read a letter from someone you respect.
6. **Take notes**
Switch to the Notes tab. Capture what resonated, what surprised you.
7. **Iterate**
Refine your prompt. Adjust your lenses. Run again.

The engine supports multiple models—you can switch between them in the model dropdown. Different models have different strengths. Some are more analytical. Some are more creative. Experiment.

Every critique is saved to History. You can copy the output, export it as Markdown, or just let it sit there, waiting for you to come back to it with fresh eyes tomorrow morning.

Deep Dives

Sometimes the surface critique points to a crack, and the only way to understand the crack is to follow it down.

The Deep Dive tab is for when the Critique Engine surfaces something interesting and you want to go further. It takes a single lens, a single topic, and a focus mode, and produces a longer, more detailed analysis.

Where critique gives you breadth—multiple lenses, a survey of the landscape—deep dives give you depth. You're drilling into one vein, following one thread, seeing how far it goes.

Focus Modes

Each deep dive lets you choose a focus mode. These shift the angle of analysis:

Strengths

What's working well through this lens? What should be preserved and amplified?

Weaknesses

Where does the design falter under this lens? What's fragile?

Opportunities

What could be done differently? Where is there room to grow?

Tensions

Where does this lens conflict with other aspects of the design? What tradeoffs are embedded?

Precedents

What existing work or theory relates to what this lens reveals?

The deep dive output is richer, more discursive, more willing to speculate. It reads less like a report and more like an essay. That's intentional. At this depth, certainty gives way to exploration.

Thoughts & Notes

The sidebar isn't just navigation. It's a journal.

The left sidebar in DesignForge holds what we call Thoughts—containers for different projects, ideas, or critique sessions. Each Thought is a named workspace that carries its own set of sources, critiques, and notes. When you switch between Thoughts, the entire studio context shifts.

Think of Thoughts as files in a cabinet. One for the app redesign. One for the brand audit. One for that personal project you're noodling on weekends. They keep your critique work organized without forcing a rigid structure.

The Notes Tab

Separate from Thoughts, the Notes tab is your in-session scratchpad. Quick observations, reactions to a critique, fragments of insight that don't yet belong anywhere. Every note is timestamped. You can search them. They persist across sessions.

Notes are intentionally minimal—plain text, no formatting, no categorization. The lack of structure is a feature. When you're mid-critique and something clicks, you don't want to spend cognitive energy deciding what folder a thought belongs in. You just want to write it down.

Capture first. Organize later. Or don't organize at all. Some of the best insights live in the margins.

History & Export

Every critique and deep dive is automatically saved to the History tab. It's a timeline of your thinking—what you asked, what the engine said, which model you used, which lenses were active. Over time, it becomes a record of how your understanding of a project evolved.

Exporting Your Work

DesignForge offers several export paths:

Markdown Export

Any critique or deep dive can be exported as a .md file. Clean, portable, version-control friendly. Drop it into your project repository, share it in a pull request, paste it into a wiki.

Full State Export

The Export button in the top bar saves your entire DesignForge state—sources, lenses, critiques, notes, settings—as a JSON file. This is your backup. It's also how you share a complete studio setup with a colleague.

Copy to Clipboard

Every output has a copy button. Quick, simple, gets the critique into whatever tool you're actually writing in.

Import and Continuity

The Import button loads a previously exported JSON state. Everything comes back—your sources, your custom lenses, your history. This makes DesignForge portable. Work on your laptop, export, import on your studio machine, and pick up exactly where you left off.

State is also automatically saved to your browser's local storage. Close the tab, come back, and your studio is waiting. But don't rely on this alone—browsers clear storage, machines crash, accidents happen. Export regularly.

Settings & Customization

The Settings tab is the engine room. Here you configure the backend, adjust visual preferences, and manage your data.


Backend Configuration

At the top of Settings, you'll find the backend selector. Four options, as we discussed in Chapter Two. Each option reveals its own configuration fields—URL and model for Ollama, API key and model for Anthropic. Changes take effect immediately; the status dot updates in real time.


Theme Toggle

The half-moon icon in the top bar switches between dark and light mode. Dark mode is optimized for extended screen work—warm blacks, low contrast, gentle on the eyes during late sessions. Light mode is optimized for readability in well-lit environments and for screenshots and screen sharing.

DARK PALETTE



LIGHT PALETTE



Vignette Control

The vignette is a subtle darkening at the edges of the screen that draws focus toward the center. It's on by default at 60%. You can adjust or disable it in Settings. Some people love it—it creates a sense of intimacy, like working inside a focused beam of light. Others find it claustrophobic. Neither reaction is wrong.

Data Management

At the bottom of Settings, you'll find the nuclear options: Export State, Import State, and Reset All. Reset clears everything—sources, lenses, history, notes, settings. It asks for confirmation first, but once confirmed, the data is gone. Use it when you want a truly fresh start, or when you're handing the studio to a colleague.

The Interface in Detail

Let's walk through every element of the interface. Not because you'll need to memorize this, but because understanding the anatomy helps you work faster when things click, and troubleshoot smarter when they don't.

The Sidebar

Fixed width, 200 pixels. At the top: the DesignForge title in small caps, a search bar for filtering your Thoughts. Below: a list of your Thoughts, each showing its name and a preview snippet. Click to switch. At the bottom: a "+ new thought" button.

The sidebar is intentionally austere. It's navigation, not decoration. On mobile, it collapses behind a hamburger menu—tap the three-line icon in the top-left to reveal it, tap anywhere else to dismiss.

The Top Bar

A thin strip across the top of the main panel. On the left: Export and Import buttons, a separator, and the theme toggle. On the right: the backend status indicator—a colored dot and a label. Green means connected. Yellow means checking. Red means something's wrong.

The top bar stays visible across all tabs. It's your always-present control strip.

The Tab Bar

Seven tabs in a horizontal row, scrollable on narrow screens. The active tab has a subtle purple underline. The tabs are labeled in uppercase monospace—Sources, Lenses, Critique, Deep Dive, Notes, History, Settings. Click to switch panels.

Each panel has a title, a description in italics, and then the panel-specific content. The description is written in *Crimson Text*, the serif font that defines DesignForge's voice—warm, literary, slightly anachronistic in a world of sans-serifs.

Input Fields

All text inputs use the same styling: a dark background (light in light mode), a subtle border, Crimson Text at 15 pixels. Textareas grow as you type. Focus is indicated by a purple border shift—not dramatic, just enough to confirm where your cursor is.

Buttons

Three button styles throughout the interface:

- Standard buttons: bordered, muted, for secondary actions like copy and export.
- Go buttons: bordered in accent purple, filled on hover. For primary actions like running a critique.
- Danger buttons: red-tinted, for destructive actions like clearing sources or resetting state.

Output Blocks

Critique and deep dive outputs appear in styled blocks—dark background, bordered, with a meta line showing the date, model, and active lenses. The output text renders in Crimson at 15 pixels with generous line height. It supports basic Markdown: headings, bold, italic, code, and lists.

Above each output: a copy button and an export button. Below: silence. The output speaks for itself.

The Toast System

Small notifications slide in from the bottom-right corner. A title, a message, a purple left border. They confirm actions: "Copied," "Exported," "State saved." They linger for three seconds, then fade. Quiet confirmations that something happened—the interface equivalent of a nod.

Keyboard Shortcuts & Tips

DesignForge has a minimal shortcut set—intentionally. The interface is designed to be discoverable through clicking and exploring. But for power users, these save time:

Ctrl/Cmd + S Save state manually. A safety net above auto-save.

Ctrl/Cmd + E Export state as JSON. Your one-key backup.

Working Tips

- Keep your source vault lean. Three strong documents beat thirty weak ones.
- Write critique prompts in second person. "Does this layout..." pulls better analysis than "Analyze the layout."
- Use deep dives on the most surprising finding from your critique, not the most obvious one.
- Export your state before experimenting with new lens configurations. It's your undo button.
- Read critique outputs the next morning. Distance improves interpretation.
- Create separate Thoughts for separate projects. Cross-contamination muddies the critique.
- Don't chase a perfect critique prompt. The first one is usually the most honest.
- Toggle the vignette off when screen-sharing. Your colleagues might not share your love of atmospheric darkness.
- When a critique feels wrong, check your sources first. The engine can only work with what it has.
- Save interesting critiques as notes. Build a personal library of design observations.

Troubleshooting

Every tool breaks eventually. The question is whether it breaks loudly or quietly.

The status dot is red

The engine can't reach your backend. If you're using Ollama, make sure the server is running. Open a terminal and type the serve command. Check that the URL in Settings matches where Ollama is actually listening. If you're using Anthropic, verify your API key is correct and your account has credits.

The status dot stays yellow

It's still checking. Give it a few seconds. If it stays yellow for more than ten seconds, the backend might be slow to respond. Try refreshing the page. If you're running a local model, it may be loading into memory—larger models take longer.

Critique output is empty or generic

Check your source vault. If it's empty or unindexed, the engine has nothing to work with. Also check your lenses—if none are active, the engine has no framework for analysis. And revisit your prompt. Vague inputs produce vague outputs.

Indexing fails

The embedding model might not be loaded. If you're using Ollama, make sure you've pulled the embedding model specified in Settings. The default is nomic-embed-text. Run the pull command in your terminal and try again.

The interface looks broken on mobile

DesignForge's mobile layout uses a collapsible sidebar. Tap the hamburger icon to access Thoughts. If the tab bar is cut off, scroll it horizontally. The interface is optimized for tablets and desktops; phones work, but it's like writing a novel on a postcard.

State disappeared after clearing browser data

DesignForge stores state in local storage. Clearing browser data erases it. This is why regular JSON exports matter. If you've lost state and have no export, the data is unrecoverable. Start fresh.

The critique contradicts itself

That might be a feature, not a bug. If your lenses have conflicting values—say, accessibility and visual density—the engine will surface that conflict. The contradiction is the insight. Read it as tension, not error.

Export produces an empty file

Make sure you've actually run a critique or deep dive before exporting. The export function pulls from the most recent output. If there's no output, there's nothing to export. Run a critique first, then export.

The model generates unrelated content

This usually means your source vault is thin or unindexed. The engine needs context to ground its analysis. Without relevant sources, it falls back on general knowledge, which feels generic. Upload more specific materials and re-index.

Performance is slow with large source files

Large files take longer to index and query. Consider splitting lengthy documents into focused sections. A 50-page research report can be broken into methodology, findings, and recommendations as separate files. Smaller, targeted documents produce faster and more relevant critique.

Design Philosophy

DesignForge was built by someone who got tired of critiquing alone.

There's a particular loneliness to solo design work. You make a hundred decisions, each one reasonable in isolation, and by the end you can't see the whole anymore. You're too close. You need distance, and distance requires another set of eyes.

DesignForge is an attempt to provide those eyes—not human eyes, not replacement eyes, but augmented eyes. A multiplier for your own critical faculty. It doesn't replace a design critique with colleagues, but it means you don't have to wait for one. You can run a critique at midnight, on a Sunday, in the middle of a creative sprint when the work is still wet.

Why Lenses?

The lens metaphor isn't decorative. It comes from a real observation about how critique works: people don't critique holistically. They critique through frameworks, whether they name them or not. One designer notices typography first. Another notices whitespace. A third sees the user journey. DesignForge makes these implicit frameworks explicit and adjustable.

Why Crimson Text?

The choice of a serif font for a design tool is deliberate. Crimson Text says: this is about reading, about language, about the slow work of understanding. It resists the tech-product impulse toward geometric sans-serifs and frictionless surfaces. It introduces a little resistance, a little texture, a reminder that critique is a literary act.

Why the Dark Interface?

Two reasons. First, practical: designers spend hours staring at screens, and a dark interface reduces eye strain during extended sessions. Second, atmospheric: the dim, muted palette creates a sense of intimacy and focus. It says: this is a studio, not a showroom. The work happens here. The light is on the work, not on the walls.

But we built light mode too, because not everyone wants to work in a cave. And that's fine.












Color & Typography Tokens

DesignForge's visual language is built on a token system—named values that define every color, font, and spacing decision. Understanding these tokens helps if you want to customize the interface or build extensions.












Typography

Body Font	Crimson Text, 15px, line-height 1.7
UI Font	IBM Plex Mono, 11px, letter-spacing 0.06em
Panel Titles	Crimson Text Bold, 22px
Section Heads	10px uppercase, letter-spacing 0.12em, accent color
Tab Labels	10px uppercase, letter-spacing 0.1em
Sidebar Items	IBM Plex Mono, 11px

Color Tokens — Dark Theme

	--bg	#0b0a0e	Primary background
	--bg2	#13121a	Card/input background
	--bg3	#1a1924	Hover states
	--tx	#d4d0de	Primary text
	--txd	#6a6578	Dimmed text
	--txg	#2e2a38	Ghost text
	--acc	#b48ead	Accent (lavender)
	--acc2	#926d8b	Accent alt
	--acc3	#d4a0c7	Accent light
	--brd	#262333	Borders
	--ok	#8aad8e	Success/connected

Color Tokens — Light Theme

	--bg	#f2ede6	Primary background
	--bg2	#e9e3da	Card/input background
	--bg3	#dfd8cd	Hover states
	--tx	#1a1520	Primary text
	--txd	#7a7088	Dimmed text
	--txg	#c8c0d0	Ghost text
	--acc	#8a5e80	Accent (mauve)
	--acc2	#6d4868	Accent alt
	--acc3	#a47098	Accent light
	--brd	#c8c0d0	Borders
	--ok	#5a8a5e	Success/connected

Spacing & Layout

Base unit	16px (--ws)
Sidebar width	200px fixed
Top bar height	40px
Tab bar height	32px
Panel padding	28px vertical, 32px horizontal
Card padding	16px
Border radius	2px (cards), 9px (toggles)
Transition	200ms ease (--tr)

Visual Effects

DesignForge layers three atmospheric effects over the interface:

- Noise Layer: Fractal noise at 4% opacity, overlay blend mode. Creates subtle texture.
- Fine Noise: A second, finer noise layer at 2% opacity, soft-light blend.
- Ambient Glow: Two radial gradients in accent colors, 6% opacity, slowly animating. Creates a sense of warmth.
- Vignette: Radial gradient darkening edges, adjustable opacity, default 60%.

Real-World Workflows

Theory is comfortable. Practice is where the blisters form.

Let's step away from features and into scenarios. Here are five workflows drawn from real use—situations where DesignForge solved a problem that no amount of squinting at a screen could fix.

Workflow 1: The Solo Brand Audit

You're a freelance designer. A client asks you to evaluate their brand system before a refresh. You have their style guide, their website copy, their mission statement, and three competitor style guides.

Upload everything into the Source Vault. Primary: client's style guide and mission. Reference: competitor guides. Scratch: your initial observations after browsing their site.

Apply the Brand Auditor preset. Add a custom lens: "Voice Consistency—Does the visual language match the written tone across touchpoints?" Run a critique with the prompt: "Where does this brand system feel coherent, and where does it fracture?"

The output becomes the backbone of your audit report. Export as Markdown, polish it in your writing tool, deliver to the client. Total time saved: roughly a day of staring and scribbling.

Workflow 2: Research Synthesis

You've finished ten user interviews. The transcripts are cleaned, each saved as a text file. You need to find patterns, contradictions, and implications for the design.

Upload all transcripts to Primary Sources. Apply the Researcher preset. Write a prompt: "What are the three most consistent pain points across these interviews, and what design implications does each one carry?"

Then deep dive into the most surprising finding. Choose the Tensions focus. The engine will show you where user needs conflict with each other—the moments where you can't please everyone and need to make a conscious tradeoff.

This isn't a replacement for reading the transcripts yourself. You've already done that. This is a second pass with a different kind of attention—one that holds all ten conversations in mind simultaneously, which is something human brains struggle with.

Workflow 3: The Late-Night Design Review

It's 11 PM. You've been working on a dashboard layout for three days. The design is due tomorrow. You know something is off, but you can't identify it. Your colleagues are asleep.

Open DesignForge. Create a new Thought called "Dashboard—Final Check." Upload your design brief and your notes about the intended user flow. Apply the Critic preset with extra weight on the Hierarchy and Readability lenses.

Prompt: "I've been staring at this layout for three days. What am I not seeing?"

The engine doesn't know what you're not seeing, obviously. But by running your brief and notes through multiple critical lenses, it might surface the disconnect you've been feeling. Maybe the data hierarchy doesn't match the task hierarchy. Maybe the primary action is buried.

Whatever it surfaces, you now have a specific thing to look at instead of a diffuse anxiety. That's worth the ten minutes.

Workflow 4: Teaching Design Critique

You're a design instructor. Your students default to vague praise or unstructured criticism. "I like the colors" isn't critique. "The blue-to-orange transition creates a temperature clash that undermines the approachable tone of the copy" is.

Set up DesignForge with a student's design brief in the vault. Walk through the lens system in class. Show how activating different lenses changes what the engine notices. Then ask students to write their own lenses and compare outputs.

The tool becomes a teaching aid: not a replacement for human critique, but a model of what structured, lens-based critique looks like in practice.

Workflow 5: The Design System Handoff

You're leaving a project. The design system you built needs to be handed off. Instead of a flat documentation dump, you create a DesignForge state file that includes your design principles, component guidelines, and custom lenses that encode your design philosophy.

The new team imports the state. They don't just get documentation—they get a critique framework tuned to the specific values of this design system. When they make changes, they can run critiques through your lenses and check consistency.

It's not perfect. No handoff is. But it's better than a Figma file and a wish.

Workflow 6: The Competitive Landscape

You're redesigning a SaaS product and need to understand how competitors solve similar problems. Collect their marketing copy, feature descriptions, and onboarding flows as text files. Upload them all as Reference Materials.

Create custom lenses: "Differentiation—Where do we offer something these competitors don't?" and "Pattern Convergence—Where do all competitors make the same choices, and what does that tell us about user expectations?"

The critique won't tell you what to build. But it will map the territory—where the field crowds together, and where there's empty space worth exploring.

Workflow 7: The Accessibility Audit

Upload your WCAG compliance checklist as a Primary Source. Add your design documentation and component specifications. Create lenses that target specific accessibility dimensions: color contrast, keyboard navigation, screen reader compatibility, cognitive load.

Run a critique asking: "Where does this design create barriers for users with disabilities, and what specific patterns should change?" Then deep dive into each barrier the engine identifies.

This doesn't replace automated accessibility testing. But it adds a layer of contextual analysis that automated tools miss—the difference between passing a contrast ratio test and actually being readable by someone with low vision in bright sunlight.

Workflow 8: Content Strategy Review

Upload your content guidelines, sample copy, and user personas. Build lenses around readability, tone consistency, information hierarchy, and emotional resonance.

Prompt: "Does the language we use match the language our users use? Where are we talking past them?"

The engine can't hear your users talk. But if you've uploaded interview transcripts alongside your marketing copy, it can spot the gap between how your users describe their problems and how your product describes its solutions. That gap is where clarity goes to die.

Advanced Techniques

Once you've spent a few weeks with DesignForge, you'll develop instincts about how to get the most from it. Here are techniques that emerge from extended use.

The Dialectical Critique

Run the same prompt twice with opposing lens configurations. First, activate only aesthetic lenses—visual harmony, typography, color emotion. Then switch to functional lenses—usability, accessibility, task efficiency. Compare the two outputs.

The gap between them is where your design lives. Every design negotiates beauty and function, and seeing both sides articulated separately makes the negotiation visible.

The Temporal Stack

Run critiques at different project stages and compare them. Week one: brief and initial concepts. Week three: refined designs and user feedback. Week six: final deliverables. The evolution of the critique mirrors the evolution of your understanding.

Over time, you build a narrative arc of how the critical landscape shifted. This is useful for post-mortems and for catching scope creep—if your week-six critique contradicts your week-one brief, something drifted.

The Outsider Lens

Create a lens called “First Impression” with the core question: “If someone encountered this design with zero context, what would they understand, feel, and do in the first five seconds?” This lens fights the curse of knowledge.

Pair it with “Expert Eye”: “What would a senior designer with twenty years of experience notice that a junior designer would miss?” The tension between naive and expert produces some of the most actionable critique.

Source Layering

Instead of uploading all sources at once, add them in layers and run critiques after each addition. Start with just the brief. Critique. Add research. Critique again. Add competitors. Each layer changes the critique in ways that reveal which inputs actually influence your design thinking.

Prompt Chaining

Use the output of one critique as the basis for your next prompt. If the first critique flags a hierarchy problem, your second prompt might be: “The previous critique identified a hierarchy issue. What specific changes to type scale, spacing, and color weight would resolve it?”

Each critique builds on the last, creating a chain of increasingly specific analysis. Three or four links usually

Collaboration & Sharing

Design is a team sport played by people who sometimes prefer to work alone. DesignForge respects both impulses.

DesignForge is primarily a solo tool. But collaboration enters through the export system.

Sharing a Critique

Export a critique as Markdown. Paste it into a shared document, a Slack channel, a pull request comment. The prose format makes it readable by non-designers—product managers, developers, stakeholders—without requiring them to understand the tool.

Sharing a Studio Setup

Export your full state as JSON. A colleague imports it. They now have your sources, your lenses, your critique history. They can build on your work or modify the lenses to reflect their own perspective.

This is powerful for onboarding new team members. Instead of explaining your critique philosophy verbally, you hand them a state file that embodies it. They learn by using.

Building a Team Lens Library

Over time, a team builds a shared library of custom lenses. An accessibility lead contributes WCAG-focused lenses. A brand strategist contributes voice-consistency lenses. A systems designer contributes component-reuse lenses.

When a new project starts, the team assembles a lens set from the library. The combination is project-specific but built from shared vocabulary. It's a way of scaling critique culture without requiring everyone in the same room.

The best teams critique like orchestras—each instrument plays a different part, but they're reading the same score.

Ethics of AI Critique

Let's talk about what DesignForge is not, because this matters as much as what it is.

DesignForge is not a designer. It doesn't have taste. It doesn't have lived experience. It doesn't know what it feels like to be confused by a navigation pattern, or delighted by a micro-animation. It processes text through statistical patterns and generates plausible critique. Sometimes that critique is genuinely insightful. Sometimes it's confident nonsense.

Your job—always—is to be the judge. Read the critique the way you'd read a review by a stranger: with interest, with openness, but without surrendering your own judgment.

Privacy Considerations

If you're using Ollama or another local backend, your data stays on your machine. Nothing leaves your network. This matters when you're working with confidential client materials or sensitive user research.

If you're using a cloud backend, your prompts and sources are sent to an external server. Don't upload materials that your client hasn't authorized for cloud processing. When in doubt, run local.

The Automation Trap

It's tempting to automate critique—to set up a lens configuration and run every design through it mechanically. Resist this. Critique is a thinking process, not a quality gate. The value isn't in the output; it's in the act of formulating the question, reading the response, and letting it change how you see.

If you find yourself skimming critiques and clicking "next," stop. You've turned a thinking tool into a checkbox. Close the tab. Go for a walk. Come back when you're ready to read slowly again.

The machine doesn't think. You do. Don't outsource that.

Bias and Blind Spots

Language models carry biases from their training data. DesignForge’s critique engine will sometimes reflect Western design conventions as universal truths, privilege certain aesthetic traditions over others, or miss cultural context that would change the analysis entirely.

If you’re designing for a culture or context you don’t share, the engine’s critique is a starting point, not an authority. Seek human reviewers from that context. No amount of lens-tuning replaces the lived experience of the people you’re designing for.

When to Ignore the Critique

Always give yourself permission to ignore the engine. If a critique says your layout breaks convention, maybe breaking convention is the point. If it flags your color palette as low-contrast, maybe the subtlety is intentional and the audience expects it.

The purpose of critique isn’t obedience. It’s illumination. A critique that you consciously reject is just as valuable as one you accept, because both sharpen your understanding of what you’re doing and why.

The Future of AI-Assisted Critique

DesignForge sits at a particular moment in the relationship between designers and AI tools. The technology will change—models will get sharper, faster, cheaper. Visual analysis will arrive. Real-time critique during the design process will become possible.

But the core question will remain the same: what does it mean to look at your own work honestly? Tools evolve. That question doesn’t. Whatever DesignForge becomes, it will always be a mirror—and what you see in the mirror depends on whether you have the courage to look.

At a Glance

A single-page reference for the moments when you need a reminder, not a chapter.

To run a critique:
Sources → upload & index → Lenses → activate → Critique → write prompt → Critique

To deep dive:
Critique first → Deep Dive → pick lens → pick focus → Go Deeper

To save your work:
Ctrl/Cmd+S for state → Ctrl/Cmd+E for export → Copy for individual critiques

To switch themes:
Half-moon icon in top bar

To add a custom lens:
Lenses tab → scroll to bottom → name + core question → Add

To share your setup:
Export → share JSON file → colleague imports

To start fresh:
Settings → Reset All (irreversible)

To change backends:
Settings → choose backend → configure → watch for green dot

To take notes:
Notes tab → type → add note → notes persist across sessions

To organize projects:
Sidebar → + new thought → name it → separate workspace

Terms Worth Knowing

Backend

The AI server that powers the critique engine. Ollama, Anthropic, WebAI, or StudioLM.

Embedding

A mathematical representation of text for semantic search. Created during indexing.

Lens

A named critique framework with a core question and adjustable intensity.

Preset Pack

A curated collection of lenses for a specific critique style.

Source Vault

The collection of uploaded documents in three weighted folders.

Thought

A named workspace in the sidebar that carries its own context.

Critique Engine

The system that reads sources through active lenses to produce analysis.

Deep Dive

A focused, single-lens analysis with a specific angle of inquiry.

Focus Mode

The angle of a deep dive: Strengths, Weaknesses, Opportunities, Tensions, Precedents.

State

The complete saved configuration, exportable as JSON.

Token

A named design value used consistently across the interface.

Vignette

The subtle darkening at screen edges that draws focus to center.

Noise Layer

Fractal noise overlaid for visual texture.

Ambient Glow

Slowly animating radial gradients for atmospheric warmth.

