

# UX through a Frontend Developer's eyes - practical UX protips for FEDs



# Why Frontend Developers own UX quality

Shipping a working interface has never been easier. Design tools, AI, component libraries, and modern scaffolding have dramatically lowered the barrier to getting something on screen.

Ideally, you work alongside a UX designer who handles research, prototyping, and interaction design. But that's not always the case. Sometimes there is no designer on the team. Sometimes the designer's bandwidth doesn't cover every screen, edge case, or state transition. In those moments, the frontend developer becomes the last line of defense for user experience - and often the person best positioned to improve it.

But a UI that *works* and a UI that *feels good* are different things. A large part of that gap lives at the implementation level: intermediate states, loading behavior, motion, feedback, layout stability, and interaction patterns. These aren't purely design decisions - they're engineering decisions that designers rarely spec in full detail. And they directly shape how an interface is experienced.

This article is a practical, hands-on collection of workflow improvements, UX patterns, and implementation-level pro tips - all from a frontend developer's perspective. No lengthy theory lectures (though a few concepts need a quick explanation to make sense). Jump in and try these techniques in your own projects to measurably improve how your interface feels.

## AI in modern frontend workflow

Before we dive into UX patterns, let's talk about the tools shaping how frontenders work today. AI-powered workflows are changing how we go from design to code - and understanding their strengths and limitations helps you know exactly where your own UX instincts need to kick in.

## Figma Developer Mode

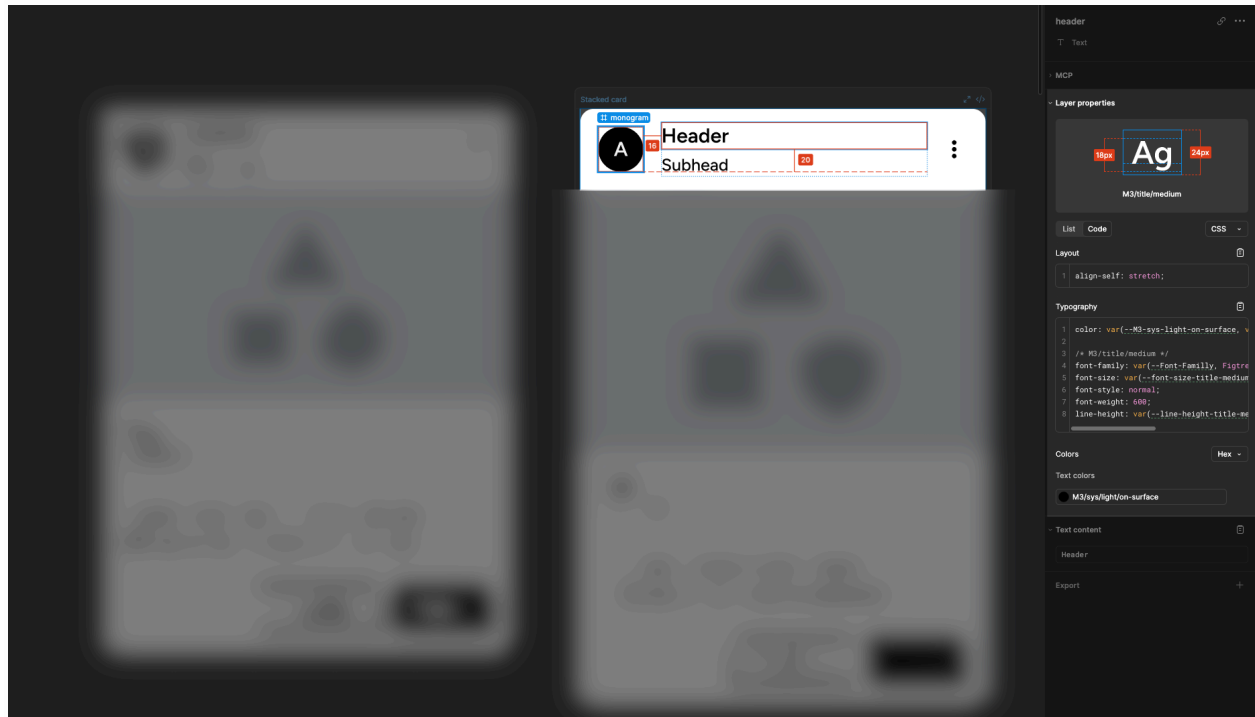
Let's start with something that significantly improves the frontend developer's life when working with a designer.

One of the most popular tools among UI/UX designers is Figma - a collaborative design platform for creating interface prototypes that eventually land on the developer's desk for implementation.

In 2023, Figma introduced a feature called `Dev Mode`, which lets developers inspect all CSS properties of any element: dimensions, margins, paddings, colors, font sizes, border radii, and much more. Think of it as "Inspect Element" but for design files - before anything is coded.



Thanks to Dev Mode, you can build pixel-perfect interfaces without guessing values or constantly pinging the designer for spacing details. It bridges the gap between what the designer intended and what you actually ship.



## AI integration with Figma and MCP

Figma also offers integration of your local AI agent - tools like Cursor or Claude Code - with an MCP (Model Context Protocol) server for Figma. MCP is an open standard (originally developed by Anthropic) for connecting AI tools to external data sources; Figma provides its own server implementation that exposes design data via this protocol. This opens up the possibility of implementing individual components or entire UIs in your chosen technology stack, based on a specific mockup, with a single prompt.

This is great for:

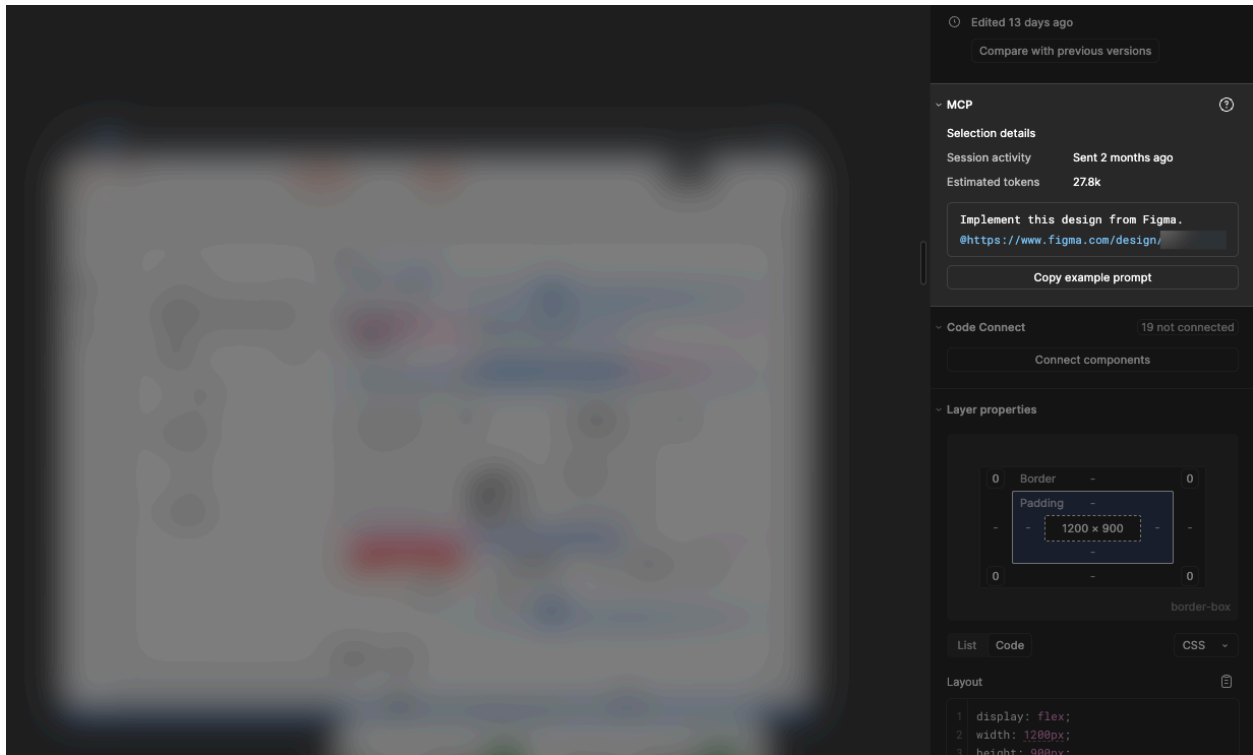
- **Rapid prototyping** - go from a Figma frame to working code in minutes, not hours. Perfect for validating ideas before committing to full implementation.
- **Exploring visual variants** - ask the AI to generate three different layouts for the same component and compare them side by side in your actual tech stack.
- **Speeding up the boring parts** - let the AI handle boilerplate styling and structural markup while you focus on the interaction logic that actually matters.



But it still requires manual refinement around:

- **Interaction states** - AI-generated code typically covers the default, “happy path” look. Hover, focus, disabled, and error states are often missing or inconsistent.
- **Dynamic data behavior** - real apps don't have three perfectly-sized cards. What happens when a title wraps to two lines? When does a list have 200 items? When a field is empty?
- **Data source integration** - AI-generated components use placeholder data. Connecting to your actual APIs, databases, or state management layer - and handling the async lifecycle that comes with it - is entirely on you.
- **Loading and error handling** - AI rarely generates skeleton loaders, optimistic updates, or graceful error boundaries. These are exactly the patterns covered later in this article.

For a step-by-step guide on setting up the Figma MCP server, check out the [official Figma MCP integration guide](#).



---

## UI Generators (Lovable, Bolt, v0)

If you don't have access to a UX designer to design your interface, you can reach for dedicated AI-powered UI generation tools.



As of early 2026, the most popular options are: [v0.dev](#), [bolt.new](#), and [lovable.dev](#). This space evolves rapidly - specific tools may change, but the strengths and limitations described below apply to AI-generated UI in general.

They all work similarly - you describe what you want in a prompt, and they generate ready-to-use interface code that you can drop into your application. It's not a perfect process; sometimes you need to regenerate the UI a few times or ask the AI to tweak individual elements. But the general rule holds: **the more specific your prompt, the better the output.**

Your prompt controls which technologies are used, though each platform has internal system prompts with sensible defaults (v0 defaults to Next.js and Tailwind, for example).

These services also support full-stack generation. For instance, v0 can create server components, API endpoints in Next.js, and fully wired-up frontend actions that connect to those backend functions - all from a single prompt.

Excellent for:

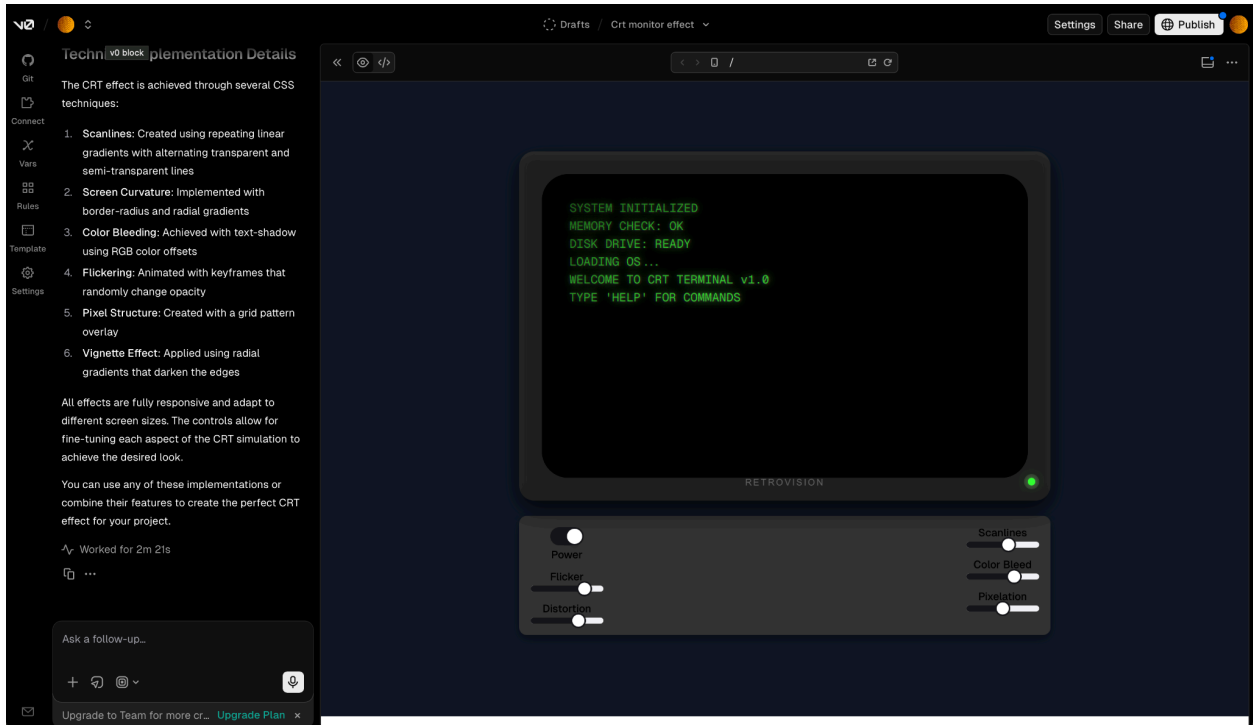
- **MVPs and proofs of concept** - when you need a working UI in hours, not days, and pixel-perfection is less important than getting something tangible in front of stakeholders.
- **Initial scaffolding** - generating a component library baseline or page layout that you then customize, rather than starting from a blank file every time.
- **Structural exploration** - trying out different page layouts, navigation patterns, or component arrangements before committing to one approach.

Commonly weak in:

- **Loading states** - most generated UIs jump straight from “nothing” to “everything rendered” with no skeleton screens, spinners, or progressive loading.
- **Error states** - what happens when an API call fails? When a form submission is rejected? Generated code rarely addresses these scenarios.
- **Async flows** - real apps deal with network latency, retries, race conditions, and stale data. AI-generated code almost never accounts for this.
- **Microinteractions** - no hover feedback, no transition animations, no subtle visual cues that make an interface feel alive.
- **Feedback patterns** - no toasts, no success confirmations, no undo prompts. The user clicks a button and has to guess whether anything happened.

Generated UI is usually a solid baseline - but it's not a finished UX. The remainder of this article covers exactly the patterns and details you need to layer on top.





## Perception and decision rules Developers should know

AI tools handle the surface-level structure well, but good UX requires understanding *why* certain layouts and interactions work. Before diving into specific implementation patterns, let's start with the foundational principles that underpin everything else.

### Gestalt Principles (practical use)

[The Gestalt principles](#) describe how the human brain organizes visual information into groups and patterns. They were formulated by psychologists in the early 20th century, but they're directly applicable to every component layout and dashboard you'll ever build.

- **Proximity groups items** - elements placed close together are perceived as related, even without lines or boxes around them. This is why spacing is one of your most powerful design tools. A form where labels are closer to their inputs than to the next field's label is instantly more readable. A dashboard where related metrics are clustered with generous space between groups communicates hierarchy without a single heading.



- **Similarity implies relation** - elements that share visual properties (color, size, shape, font weight) are perceived as belonging to the same category. This is why all your primary action buttons should look the same, why all destructive actions should share a color, and why breaking visual consistency unintentionally can confuse users.
- **Continuity guides scanning** - the eye follows smooth lines and curves. Aligned elements create visual flow. This is why left-aligned form labels are easier to scan than centered ones, why consistent column alignment in tables matters, and why a misaligned element feels “wrong” even when the user can’t articulate why.
- **Figure–ground builds layers** - the brain separates foreground content from background surfaces. This is the principle behind cards, modals, and overlays. A card with a subtle shadow “lifts” off the background and becomes the focus. A modal with a dark backdrop makes the background recede. Using lightness and shadow to create depth helps users understand what’s interactive and what’s context.

These aren’t abstract concepts - they’re the reason good layouts “just work”, and bad ones feel confusing.



**PROXIMITY**

**BAD — LABEL FAR FROM ITS INPUT**

Email

Password

**GOOD — LABEL TIGHT TO ITS INPUT**

Email

Password

---

**SIMILARITY**

**BAD — INCONSISTENT BUTTON STYLES**

**GOOD — CONSISTENT HIERARCHY**

---

**CONTINUITY**

**BAD — RAGGED LABEL ALIGNMENT**

Name

Email

Phone

**GOOD — ALIGNED LABELS**

Name

Email

Phone

---

**FIGURE-GROUND**

**BAD — NO VISUAL SEPARATION**

Order Summary  
Total: \$49.99 Shipping: Free Estimated delivery: 3-5 days

**GOOD — CARD WITH SHADOW**

Order Summary  
Total: \$49.99  
Shipping: Free  
Estimated delivery: 3-5 days

## Hick's Law

Hick's Law states that the time it takes to make a decision increases logarithmically with the number of choices. In plain terms, **more options mean slower, harder decisions.**

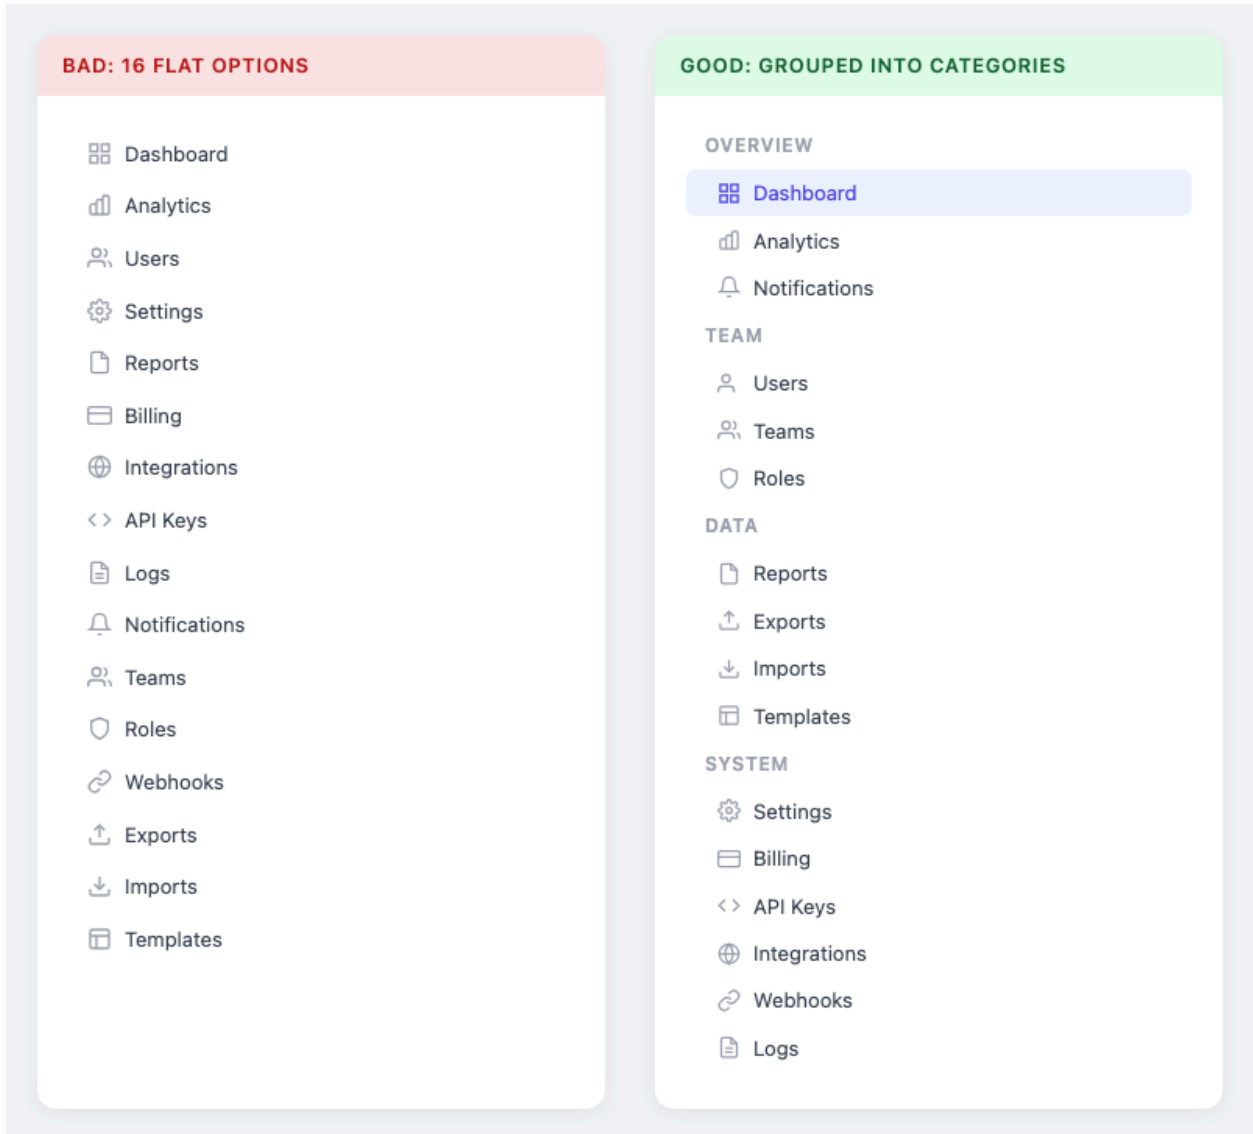


This has direct implications for navigation menus, settings pages, form dropdowns, and any interface where the user must choose from a set of options. A sidebar with 18 flat, unsorted links is measurably harder to navigate than one with 4 grouped categories that expand to reveal specific options.

Practical applications:

- **Group and categorize** - instead of a flat list of 15 options, organize them into 3–5 logical groups. The user first picks a category (easy, few choices), then picks an item within it (easy, few choices). Two simple decisions are faster than one complex one.
- **Progressive disclosure** - show essential options first, and hide advanced ones behind an “Advanced” toggle or submenu. Most users never need the advanced settings; don’t slow them down with options they’ll never use.
- **Smart defaults** - pre-select the most common option. This reduces the decision from “pick one of these” to “accept this or change it” - a much simpler cognitive task.
- **Search as an escape hatch** - for very large option sets (country selectors, font pickers), provide a search/filter input so users can jump directly to what they want instead of scanning a massive list.





## Fitts's Law

Fitts's Law predicts that the time to reach a target depends on two things: the distance to the target and the size of the target. Bigger, closer targets are faster and easier to hit.

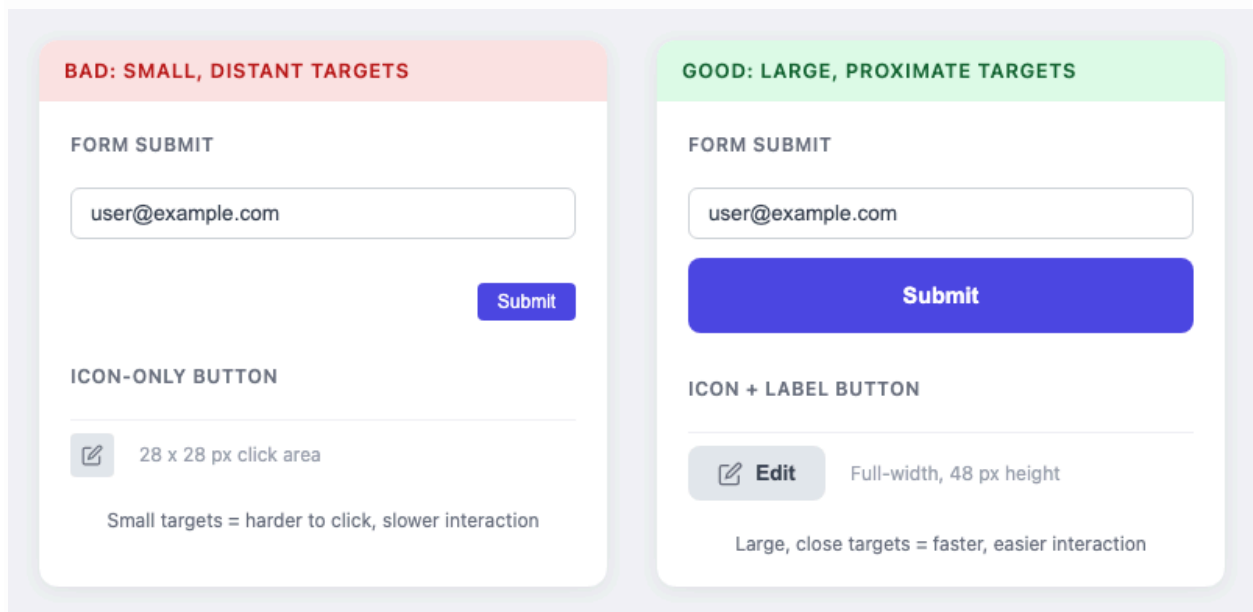
This sounds obvious, but it's violated constantly in real interfaces:

- **Tiny action buttons** - a 24x24px icon button with no padding is technically clickable, but it's a frustrating target, especially on mobile. A 44x44px minimum touch target (Apple's recommendation) or 48x48px (Google's) makes a huge difference.



- **Submit buttons far from input fields** - if a user just finished typing in a text field, their cursor (or thumb) is near that field. A submit button placed far away forces unnecessary travel. Place primary actions adjacent to the inputs they act on.
- **Icon-only buttons without labels** - a small gear icon for settings is harder to click than a button that says “Settings” with a gear icon. The text dramatically increases the clickable area and, as a bonus, removes the ambiguity of icon-only actions.
- **Mobile edge targets** - buttons placed at the corners or top of a mobile screen are the hardest to reach with one hand. Critical actions should be in the lower-center area, within the thumb’s natural arc.

The takeaway: make important things big, make them close to where the user’s attention already is, and never rely on tiny click targets for critical actions.



## Progressive disclosure

Progressive disclosure is the principle of showing only what’s necessary at each stage of an interaction and revealing more as the user requests it or it becomes relevant.

The goal is to reduce initial overwhelm and let users engage at their own pace:

- **Multi-step wizards** - instead of a single form with 30 fields, break it into logical steps (account info → preferences → review). Each step shows only a handful of fields, making the task feel manageable. Progress indicators (“Step 2 of 4”) add motivation and context.



- **Collapsible accordions** - FAQ pages, settings panels, and long documentation benefit from accordions that show section titles with content hidden until clicked. The user scans the titles, opens only what's relevant, and isn't overwhelmed by a wall of text.
- **"Advanced" and "More options" sections** - in forms, dashboards, and configuration screens, hide rarely used options behind a toggle or an expandable section. Power users know to look for it; casual users aren't distracted by options they don't understand or need.

Progressive disclosure works hand-in-hand with Hick's Law - by revealing fewer options at each step, you're reducing decision complexity at every stage.

---



### 1. MULTI-STEP WIZARDS

#### BAD: ALL 12 FIELDS AT ONCE

Name  
Full name

Email  
you@example.com

Password  
Min. 8 characters

Phone  
+1 (555) 000-0000

Address  
123 Main St

City

State

Zip  
00000

Country  
United States

Company  
Company name

Role  
Your role

Bio  
About you...

Submit

#### GOOD: BROKEN INTO STEPS

Step 1 of 3 — Account Info  
We'll start with the basics

Name  
Full name

Email  
you@example.com

Password  
Min. 8 characters

Back Continue

### 2. COLLAPSIBLE ACCORDIONS

#### BAD: WALL OF TEXT

**How do I reset my password?**  
Go to Settings → Security → Change password. Enter your current password, then type a new one. Click Save. You'll receive a confirmation email.

**Can I change my username?**  
Yes. Navigate to Settings → Profile → Edit. Your username must be unique and between 3–20 characters. Changes take effect immediately.

**How do I cancel my subscription?**  
Go to Settings → Billing → Manage plan. Click "Cancel subscription." You'll retain access until the end of your billing period.

**What payment methods do you accept?**  
We accept Visa, Mastercard, American Express, and PayPal. You can add or change payment methods in Settings → Billing.

**How do I export my data?**  
Go to Settings → Privacy → Export data. Choose the format (JSON or CSV) and click Export. You'll receive a download link via email within 24 hours.

#### GOOD: EXPANDABLE SECTIONS

**How do I reset my password?** >  
Go to Settings → Security → Change password. Enter your current password, then type a new one. Click Save.

**Can I change my username?** >

**How do I cancel my subscription?** >

**What payment methods do you accept?** >

**How do I export my data?** >

### 3. "ADVANCED" OPTIONS TOGGLE

#### BAD: ALL OPTIONS EXPOSED

**Email notifications**  
Receive updates via email

**Push notifications**  
Browser push alerts

**Webhook URL**  
POST events to custom endpoint

**Digest frequency**  
Batch notification interval

**Rate limiting**  
Max notifications per hour

**Custom headers**  
Add headers to webhook requests

#### GOOD: ESSENTIAL FIRST, ADVANCED ON DEMAND

**Email notifications**  
Receive updates via email

**Push notifications**  
Browser push alerts

Advanced options

POWER USER SETTINGS

**Webhook URL**  
POST events to custom endpoint

**Digest frequency**  
Batch notification interval

**Rate limiting**  
Max notifications per hour

**Custom headers**  
Add headers to webhook requests



# Visual Clarity: typography, color, and states

Those principles are useful in theory - now let's apply them concretely, starting with the most fundamental visual building blocks: typography, color, and component states. When you're building the interface, the decisions about font rendering, color consistency, and component state coverage are yours to make. These details compound - a consistent type scale, a cohesive color palette, and fully defined component states are what separate a "developer-built" UI from a "designed" one.

## Typography

Typography is the backbone of any interface. Users spend most of their time reading - labels, buttons, headings, paragraphs, tooltips - so the choices you make about fonts directly impact usability.

- **Stick to 1–2 font families maximum** - one for headings and one for body text is a classic, safe combination. Using more than two creates visual fragmentation - the page feels unfocused, like multiple designs were merged together. If you're choosing a single font, pick a versatile sans-serif (Inter, Open Sans, or the system font stack). Avoid defaulting to serif system fonts like Times New Roman in app UIs - they immediately signal "unstyled" or "unfinished" to modern users accustomed to clean sans-serif interfaces.
- **Limit weight variants** - within your chosen fonts, use 2–3 weights consistently: regular (400) for body text, medium (500) for emphasis, and bold (700) for headings. Using every weight from 100 to 900 creates subtle inconsistencies that make the hierarchy feel muddy.
- **Establish a consistent type scale** - pick a scale (e.g., 12px, 14px, 16px, 20px, 24px, 32px) and use only those sizes. Random sizes like 13px, 17px, and 22px scattered throughout the app create an unconscious sense of disorder. Most design systems use a mathematical scale (each step is 1.25× or 1.33× the previous).
- **Use a readable line-height** - for body text, a line-height of 1.5–1.7 relative to font size ensures comfortable reading. For headings (which are larger and shorter), 1.2–1.3 is typically sufficient. Cramped text (line-height: 1) is visually dense and tiring to read; overly loose text (line-height: 2+) feels disconnected and wastes space.



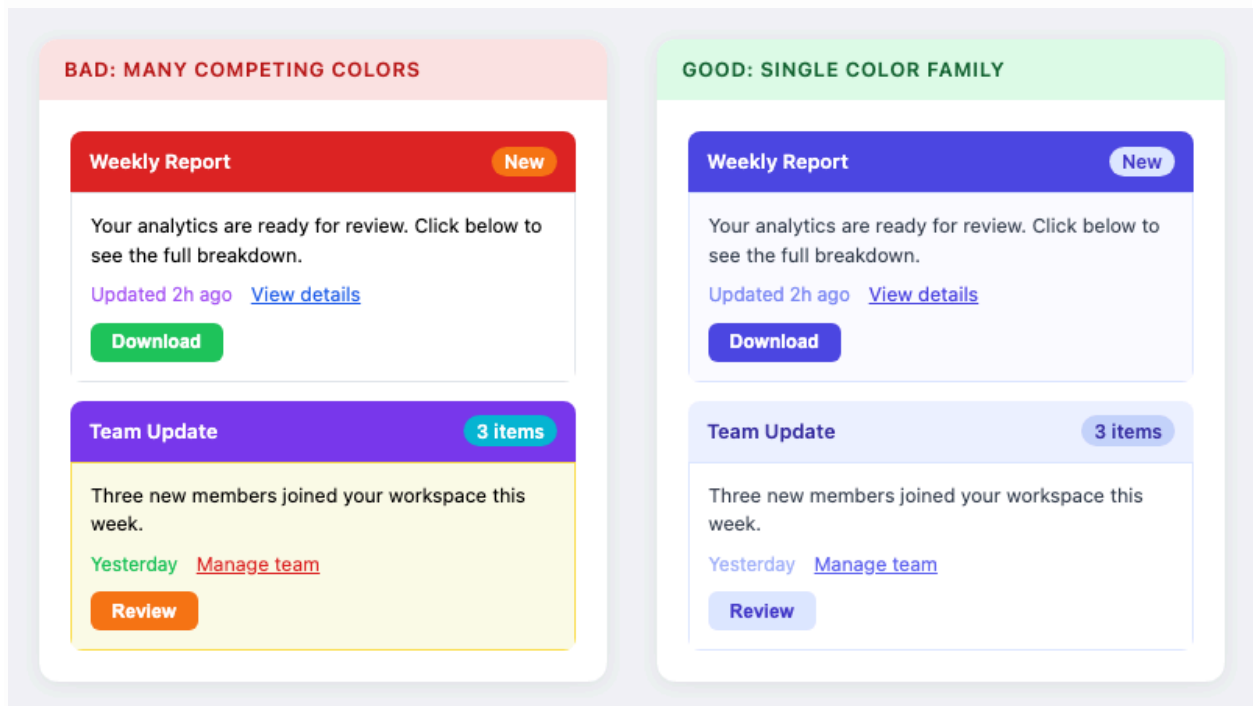
## Color and layering

A common mistake in developer-built interfaces is using too many different colors. A red header, green buttons, blue links, orange badges, and purple highlights might each seem reasonable in isolation, but together they create visual chaos.

Instead of reaching for a new color every time you need to distinguish something:

- **Use shades and tints of one base color** - pick a single brand color (e.g., indigo) and derive your entire palette from it. Light tints for backgrounds, medium shades for borders and secondary elements, the full-saturation color for primary actions, and dark shades for text. This creates instant visual cohesion. Tools like [Coolors](#) or [Realtime Colors](#) can help you generate and preview a cohesive palette in seconds.
- **Create depth through lightness, not hue** - differentiate UI layers (page background → card → elevated card → modal) by increasing lightness from dark to light (or vice versa for dark themes). This uses the Gestalt figure-ground principle to create hierarchy without introducing new colors.
- **Keep the palette consistent and meaningful** - reserve specific colors for specific meanings: red for errors and destructive actions, green for success, amber for warnings, and your brand color for primary actions. When color has a consistent meaning, users learn the visual language and navigate faster.

Color should signal state, not just decorate. Every color choice should answer the question: “What does this tell the user?”



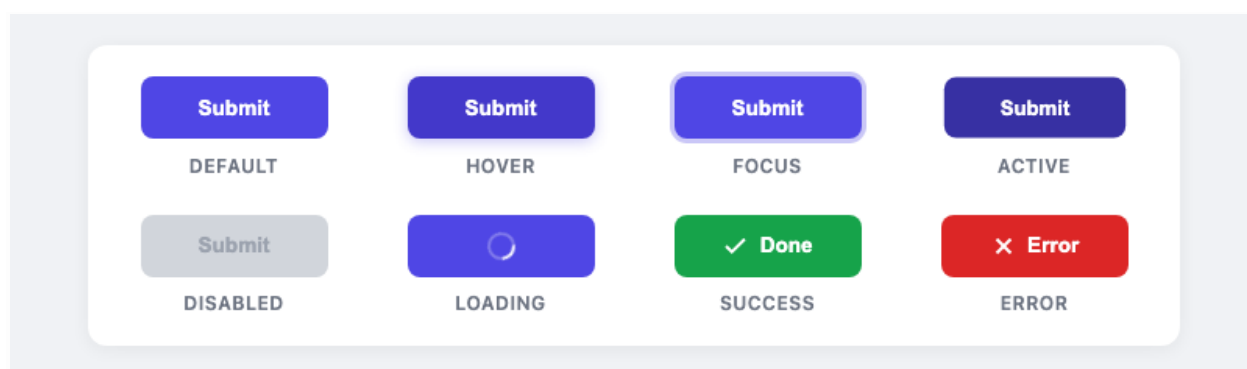
---

## Complete component states

A button that just sits there looking the same whether it's clickable, hovered, focused, loading, disabled, or errored isn't a component - it's a rectangle with text in it.

Every interactive component should define these states:

- **Default** - the resting state. What the component looks like when no interaction is happening.
- **Hover** - visual feedback on mouse-over. A subtle background change, a slight lift shadow, or a color shift that says "this is interactive, and I know your cursor is here."
- **Focus** - a visible focus ring or outline for keyboard navigation. Don't remove `:focus` styles without providing a visible alternative (more on this in the Accessibility section).
- **Active / Pressed** - a momentary visual change when the element is being clicked or tapped. Usually, or darker background. Confirms "your click registered."
- **Disabled** - grayed out, with `cursor: not-allowed` and the `disabled` attribute (or `aria-disabled` when you still want the element to remain focusable - see the Accessibility section for when and why). Note: avoid combining `cursor: not-allowed` with `pointer-events: none` - the latter prevents all pointer events, which means the cursor style won't be visible. Use one approach or the other. A common mistake is leaving a button looking clickable when it shouldn't be - for instance, a "Submit" button that should be disabled until all required fields are filled.
- **Loading** - a spinner or loading animation replacing the button text (or appearing next to it). Prevents duplicate submissions and informs the user that their action is being processed. The button should also be functionally disabled during loading.
- **Error** - red borders, tinted background, or an icon indicating something went wrong with this specific component. An input field that failed validation, a button whose action failed.
- **Success** - a brief green tint, checkmark icon, or "Done!" text that confirms a successful action before returning to the default state.



AI generators and rapid prototyping tools almost always skip most of these states. They generate the default look and call it done. It's the developer's responsibility to ensure every interactive element feels complete and responsive across all these interaction scenarios.

---

## Empty states as UX opportunities

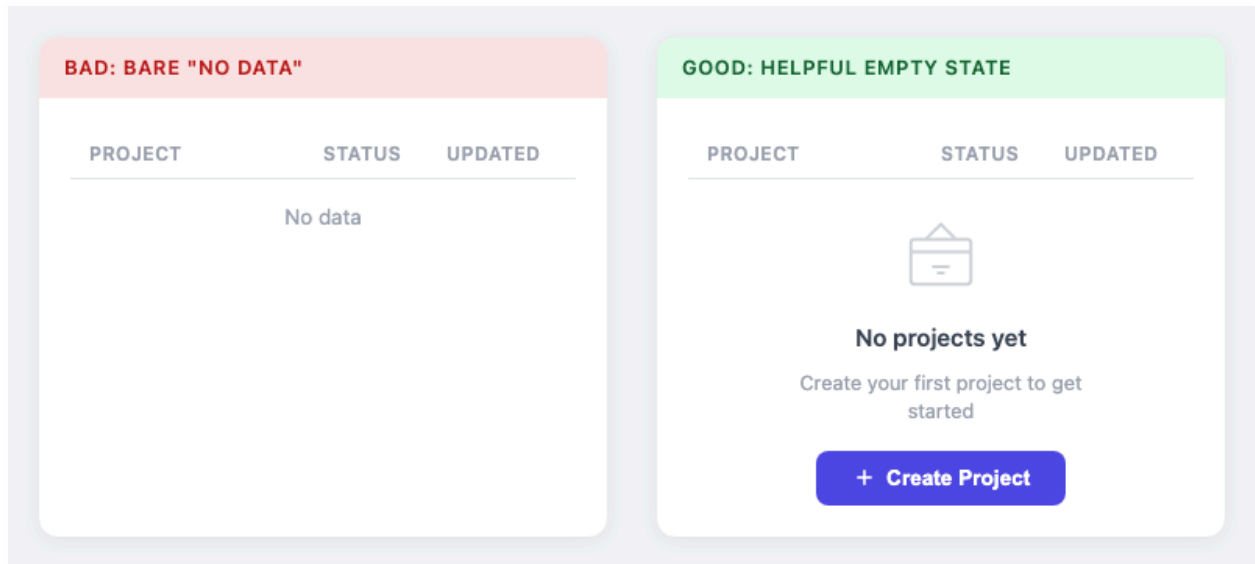
An empty state - a page, list, or section with no data yet - is one of the most overlooked UX opportunities. The default developer instinct is to show "No data" or an empty table with headers but no rows. This is technically correct and deeply unhelpful.

An empty state is the first thing a new user sees when they open your app. It's also what a user sees after clearing a list, archiving all items, or applying filters that return zero results. In every case, showing "No data" leaves the user stranded - they see nothing and don't know what to do next.

Instead, treat empty states as a chance to guide behavior:

- **Suggest the next action** - "No projects yet" is only half the message. The full message is "No projects yet - click 'Create Project' to get started." Tell the user what to do, not just what's missing.
- **Provide a hint or explanation** - briefly explain what this section will contain when populated. "Your recent activity will appear here" sets expectations and reassures the user that the feature works.
- **Include a prominent CTA** - a clear, visible "Create your first..." button right in the empty state area. Don't make the user hunt for the "New" button in a toolbar or navigation menu.
- **Add a friendly illustration** - a simple, lightweight graphic (an empty inbox icon, a folder with a plus sign) makes the empty state feel intentional rather than broken. It signals that this is a designed experience, not a missing feature.





Empty space can - and should - guide behavior rather than just occupy pixels.

---

## Performance as UX (Perceived Performance)

Your interface may look right in a static mockup - but users don't experience static pages. They see content appear, load, and shift. Perceived performance often matters more than raw performance metrics. A page that loads in 800ms but shows a blank screen feels slower than one that takes 1.2 seconds but shows skeleton placeholders from the first frame.

Some of the tips below might sound counterintuitive at first - I've even seen memes mocking the idea of a "minimum loader duration." But trust the process. These patterns exist because they've been validated across countless products, and your users will thank you for implementing them.

---

## Skeletons instead of empty screens

Nobody likes waiting. But when a user has to wait, the least you can do is let them know the page is still alive and working.



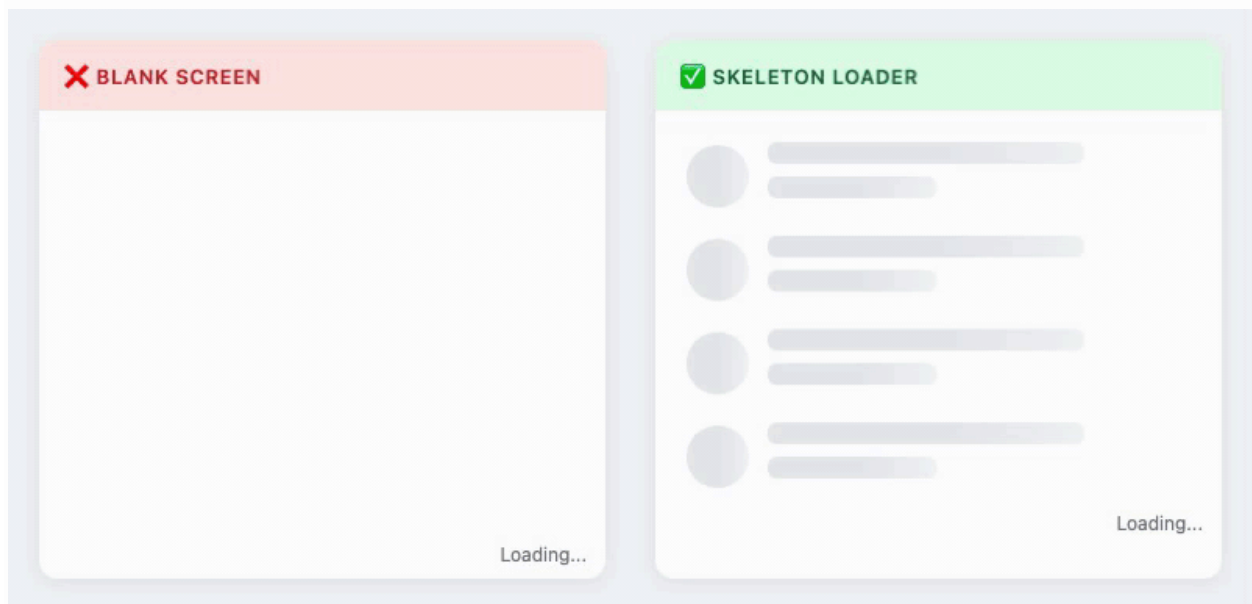
A blank screen immediately triggers the “something went wrong” alarm in a user’s brain. They start wondering: Did the page crash? Is my connection down? Should I refresh? That moment of uncertainty is a UX failure - even if the data arrives half a second later.

The current standard for handling this is to use skeleton screens - lightweight placeholder layouts that mirror the general structure of the content being loaded, paired with a subtle shimmer animation that signals “content is on its way.”

Skeleton screens beat blank pages and isolated spinners because they:

- **Preserve the page structure** - the user immediately understands the layout they’re about to interact with, even before the real data arrives. There’s no jarring rearrangement when content finally loads.
- **Preview the content shape** - rectangle blocks hint at where images, text blocks, and buttons will appear. This primes the user’s brain, making the transition to real content feel faster.
- **Reduce uncertainty and perceived wait time** - research consistently shows that users perceive skeleton-loaded pages as faster than spinner-loaded ones, even when the actual load time is identical.

They communicate progress from the very first frame - and that’s what makes the difference.



## Minimum loader duration



A loader that flashes for 30 milliseconds and vanishes looks like a rendering glitch. Users see a flicker - something appeared and disappeared before their brain could even register what it was. Instead of feeling fast, it feels broken.

A short minimum display time - typically 300 to 600 milliseconds - gives the user's brain enough time to register the loading state and understand the transition. The skeleton appeared, and the content replaced it. That's a story the brain can follow. Without that minimum duration, it's just visual noise.


This is not about artificially slowing down your system. If the data arrives in 20ms, you're only adding a couple of hundred milliseconds - imperceptible in terms of actual productivity, but a massive improvement in how polished the interface feels. The goal is to avoid the "did that just glitch?" moment and replace it with a smooth, intentional transition.

A practical implementation is straightforward: track both the fetch start time and the response time, then delay the content reveal by whichever is larger - the actual load time or your minimum duration.

```
const MIN_DURATION = 400;

const [data] = await Promise.all([
  fetchData(),
  new Promise((r) => setTimeout(r, MIN_DURATION)),
]);

setData(data);
```



## Loading animations should not be uniform

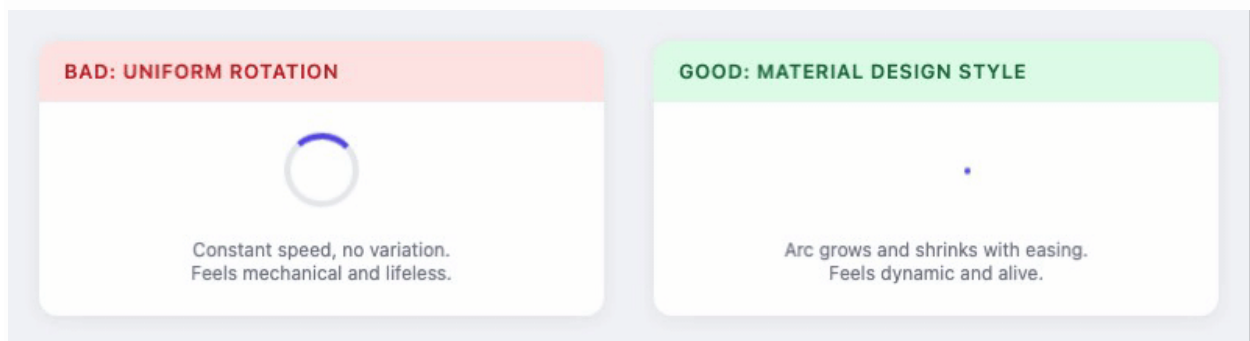
A spinner that rotates at a constant speed with a fixed arc length feels mechanical, lifeless, and - counterintuitively - slower. Our brains are wired to notice variation and interpret it as activity. Constant, uniform motion reads as "stuck" or "looping endlessly."



Better loading animations use:

- **Non-linear easing** - the animation accelerates and decelerates rather than moving at a constant pace. This mimics real-world physics, where objects don't start and stop instantly.
- **Varying arc or segment lengths** - the spinning element grows and shrinks as it rotates, creating a sense of "effort" and progress. The changing shape keeps the user's attention and suggests that something is actively happening.
- **Speed variation** - brief moments of faster and slower movement make the animation feel organic rather than robotic. Even subtle changes in velocity create a sense of life.

Google's Material Design indeterminate loader is the gold standard here. Watch how the arc stretches, contracts, and shifts speed as it circles - it feels like it's actually *doing* something, not just spinning in place.



## Progress bars when duration is unknown

When you don't know how long something will take - data fetching, file uploads, background processing - a standard linear progress bar creates a problem. If it moves at a constant speed and the task takes longer than expected, the bar either stops dead (making the user think it's frozen) or you're forced to guess the total duration (making the bar inaccurate).

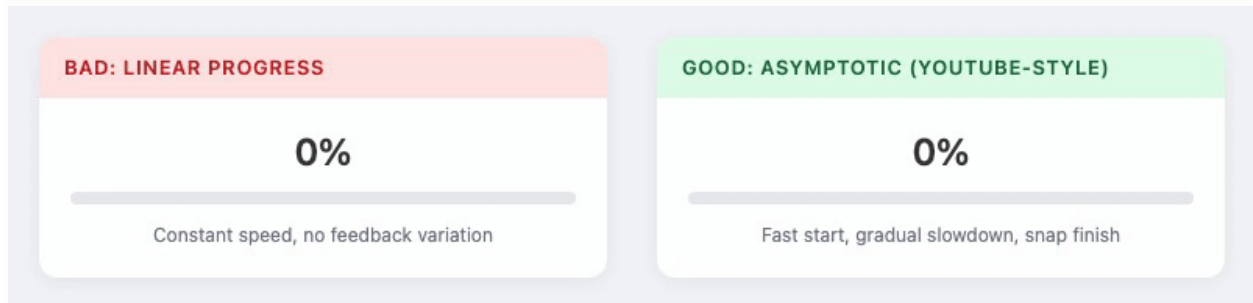
The solution is an **asymptotic progress pattern**, famously used by YouTube's top-of-page loader:

- **Fast initial progress to ~40–60%** - the bar jumps forward quickly, immediately reassuring the user that something is happening. This leverages the psychological principle that seeing early, rapid progress makes the entire wait feel shorter.
- **Gradual deceleration through the middle range** - progress slows as it approaches higher percentages, stretching the perceived timeline without ever stopping.



- **Approaching but never quite reaching 100%** - the bar creeps asymptotically toward completion, so it's always moving. When the actual task finishes, the bar snaps to 100% and completes.

Users see continuous, honest-feeling progress without false precision about timing. The bar never lies - it just optimistically front-loads the visual progress.



## Preventing layout shift (CLS as a UX issue)

Layout shifts - when elements on the page suddenly jump to a new position - are one of the most disorienting experiences a user can encounter. You're reading a paragraph, and suddenly it leaps 200 pixels down because an image loaded above it. You're about to click a button, and it moves because a banner appears. It's jarring, it breaks trust, and it's entirely preventable.

This is so impactful that Google made it one of the three Core Web Vitals: Cumulative Layout Shift (CLS). But beyond SEO implications, it's fundamentally a UX issue.

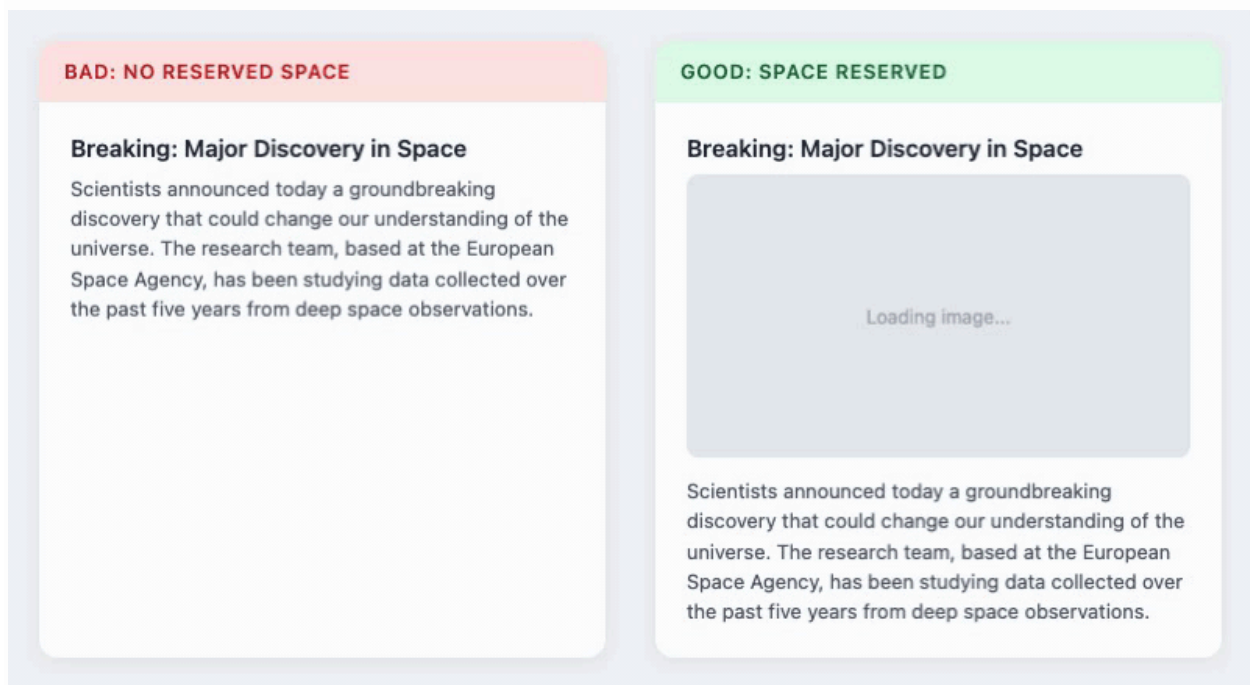
Strategies to prevent layout shifts:

- **Reserve space for dynamic content** - if you know an image is coming, set explicit `width` and `height` attributes (or use the CSS `aspect-ratio` property) so the browser can allocate space before the image downloads.
- **Use min-height on containers** - for sections that load asynchronously (ads, embeds, lazy-loaded components), set a `min-height` that approximates the expected content height. Overestimating slightly is better than not reserving any space at all.
- **Show placeholders and skeletons** - skeleton screens aren't just for perceived performance. They also lock in the layout dimensions, preventing content below from shifting when real content loads.



- **Choose the right hiding strategy** - `visibility: hidden` keeps an element in the document flow (occupying its space) while making it invisible, preventing layout shift. `display: none` removes it from the flow entirely, causing everything around it to collapse and then re-expand when it appears. Use `visibility: hidden` when you want to reserve space for content that will appear shortly. Note that both approaches also hide content from screen readers - see the Accessibility section for how to hide elements visually while keeping them accessible.
- **Be careful with dynamically injected content** - inserting elements above the user's current scroll position (cookie banners, notification bars, lazy-loaded items in a feed) almost always causes a visible shift. Insert below the viewport, or push content down with a smooth animation rather than an instant jump.

Visual stability builds trust. When the layout stays put, users feel in control - they can read, scan, and click without second-guessing where things will end up.



## UX of data operations and user actions

The page loaded smoothly - now the user starts interacting. What happens when they click a button, submit a form, or delete an item? The way your interface responds to user actions is where “functional” and “delightful” diverge the most. A half-second delay between clicking “Like” and seeing the heart fill in can make your entire app feel sluggish, even if everything else is fast.



## Optimistic update

Optimistic updates flip the traditional request flow on its head. Instead of the classic sequence - user clicks, UI shows a spinner, backend responds, UI updates - you update the UI *immediately* and assume the server will confirm. The request still happens in the background, but the user doesn't have to wait for it.

This pattern works beautifully for:

- **Likes and reactions** - the heart fills in the moment you tap it. Waiting for a round trip to the server would make the interaction feel unresponsive and broken.
- **Toggles and switches** - enabling dark mode, toggling a notification setting, or switching a filter should feel instant. The UI state change is the feedback.
- **Inline renames** - renaming a file, a project, or a label should reflect immediately in the UI. The user sees their new name right away.
- **Drag-and-drop reordering** - when a user drags an item to a new position, it should remain there immediately. Reverting to the original order while waiting for a server response would feel broken.
- **List removals** - removing an item from a list should cause it to disappear (ideally with a smooth animation) right when the user clicks remove.

The benefits are significant:

- **Instant feedback** - the UI responds in a frame. No spinners, no loading states, no perceptible delay.
- **Smoother interaction flow** - the user can perform multiple actions in rapid succession without waiting for each one to resolve.
- **Dramatically improved perceived speed** - the app feels fast even over slow connections because the UI never blocks on network requests.

A simplified optimistic toggle:

```
function toggleLike(postId: string) {  
  
  const prev = isLiked;  
  
  setIsLiked(!prev); // update UI immediately  
  
  api.toggleLike(postId).catch(() => {  
  
    setIsLiked(prev); // rollback on failure  
  
    toast.error("Couldn't update - reverted");  
  
  });  
}
```

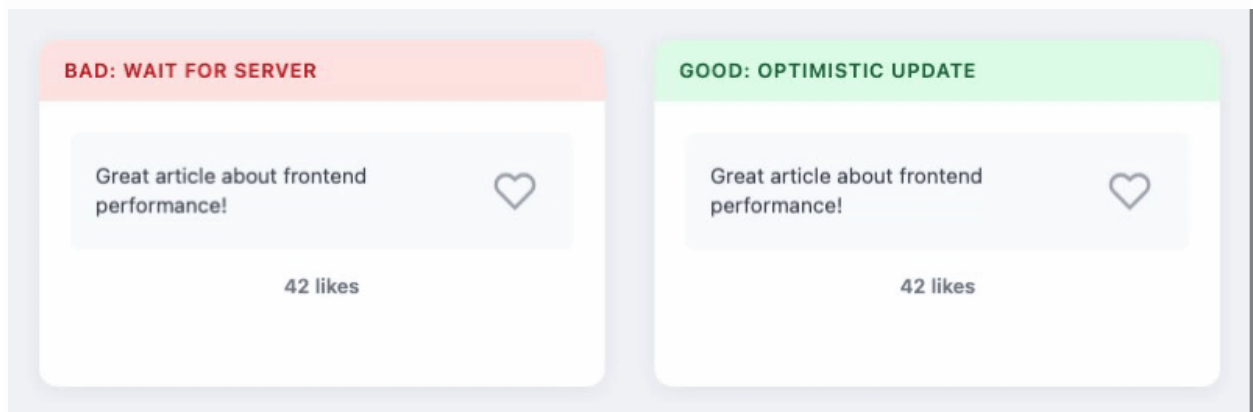


}

But optimistic updates come with engineering requirements:

- **Rollback support** - if the server rejects the change (validation error, permission issue, or conflict), you need to gracefully revert the UI to its previous state. This means keeping a reference to the pre-update state.
- **Error handling with user communication** - when a rollback occurs, show a clear, non-disruptive message explaining what happened. A toast saying “Couldn’t save - reverted” is usually enough.
- **Retry strategy** - for transient failures (network hiccup, server timeout), you may want to retry automatically before reverting.
- **Sync indicators for long operations** - for actions where the server-side effect matters (sending a message, publishing a post), consider a subtle “syncing” indicator that confirms when the backend has caught up.

One important boundary: avoid optimistic updates for high-risk or irreversible operations. Deleting an account, transferring money, or sending an email should wait for server confirmation before reflecting in the UI.



## Delete + Undo instead of confirm modals

The classic deletion pattern goes like this: user clicks delete, a modal pops up asking “Are you sure?”, the user clicks “Yes” (or more often just mashes Enter because the modal is expected and annoying), and the item gets deleted.

The problems with this flow:



- **It breaks the user's flow** - a modal is a full-stop interruption. The user has to context-switch from "I'm managing my list" to "I'm reading and responding to a dialog box." This cognitive interruption adds up fast when you're doing bulk operations.
- **It adds friction to every single deletion** - even when the user is 100% certain they want to delete something, they're forced through an extra step. Over time, this friction erodes trust in the interface's efficiency.
- **Users auto-click through it** - after seeing the same confirmation modal 20 times, the "Yes, delete" button becomes muscle memory. The modal stops being a safety net and becomes an empty ritual. When a user actually *does* need the safety net, they'll click through it anyway.

The better pattern - used by Gmail, Slack, Notion, and most modern productivity tools:

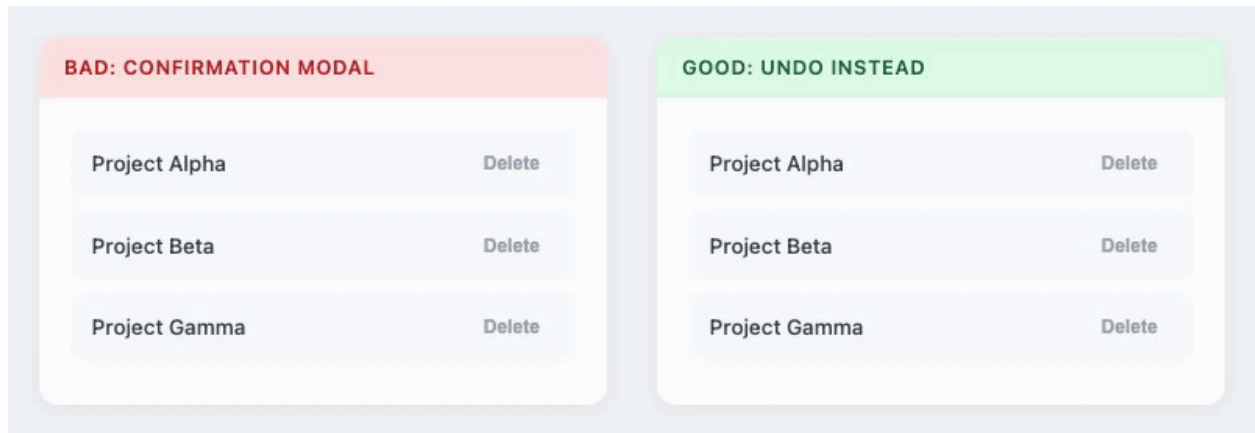
- **Remove the item from the UI immediately** - the delete happens visually right when the user clicks. No dialog, no interruption.
- **Show a toast notification with an Undo option** - a non-blocking toast slides in: "Item deleted - Undo." The user can continue working or click Undo within a few seconds.
- **Delay the actual server request** (or use soft delete) - the backend deletion happens after a short delay (5–10 seconds), giving the Undo window time to work. Alternatively, use soft delete on the server side (mark as deleted, hard-delete later) - though keep in mind that soft delete adds backend complexity: queries need to filter out deleted records, indexes grow larger, and you'll need a cleanup strategy that respects data retention policies like GDPR.

The advantages are clear:

- **No blocking dialog** - the user's flow is uninterrupted. They can keep working immediately after clicking delete.
- **Faster interaction** - one click instead of two (or three, if the modal includes a "Are you really sure?" checkbox).
- **Real, usable recovery** - unlike a confirmation modal (which users click through blindly), an Undo toast is actually useful. It catches genuine mistakes because it appears *after* the action, when the user can see the consequence and decide if it was correct.

This pattern connects naturally with the toast-based feedback discussed in the next section.





## Microinteractions and attention guidance

Actions need responses - and the quality of those responses is what separates a functional interface from a polished one. Microinteractions - small, purposeful animations and feedback loops - are how your interface communicates with the user at a subliminal level. They say “this worked,” “look here,” or “this is changing” without requiring the user to read a single word.

### Subtle motion instead of UI jumps

When something changes on screen - a panel opens, a list item appears, a section expands - the user’s brain needs a moment to understand what happened. If the change is instant (the element is suddenly *there*), the brain has to work to figure out the new state. But if the change is animated, even briefly, the motion itself tells the story.

Use small transitions instead of abrupt changes:

- **Fade** - new elements appearing with a quick opacity transition (150–250ms) feel intentional rather than jarring. This works well for content that loads asynchronously, tooltip appearances, and notification popups.
- **Scale** - a subtle scale-up from 95% to 100% when an element appears gives it a sense of “arriving.” Combined with a fade, this is the bread-and-butter entrance animation used by virtually every polished UI.
- **Slide** - elements that enter from a logical direction (a sidebar sliding in from the left, a dropdown sliding down from its trigger) help the user understand where the element “came from” and where it will “go back to” when dismissed.

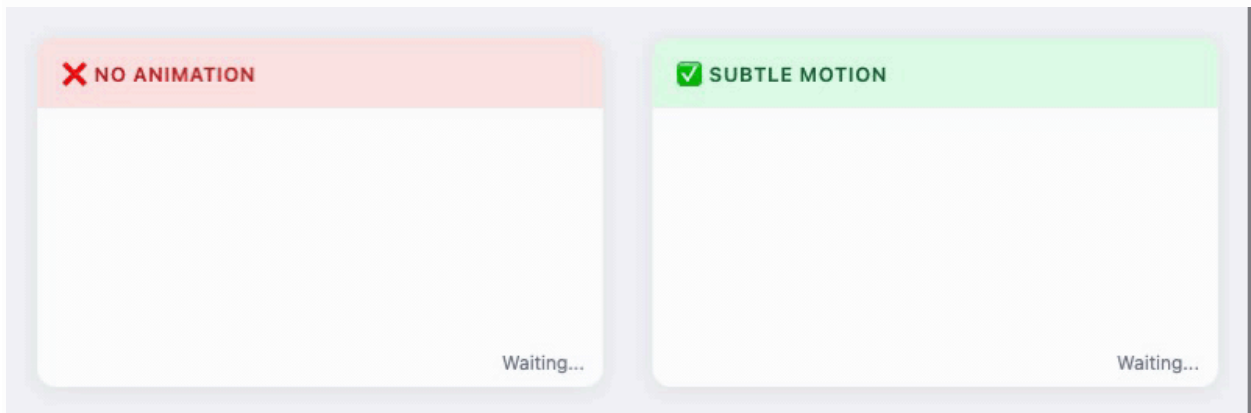


- **Easing curves** - never use linear timing for UI transitions. `ease-out` for entrances (fast start, gentle landing), `ease-in` for exits (gentle start, fast departure), and `ease-in-out` for state changes. Linear motion feels robotic; eased motion feels natural.

Interfaces feel smoother and more predictable when changes are animated. The key is keeping animations short (150–300ms) - long enough to register, short enough not to feel sluggish.

One important caveat: always respect the `prefers-reduced-motion` media query. Some users have motion sensitivity or vestibular disorders, and animations that feel polished to most people can cause discomfort or nausea for them. Provide instant (or heavily reduced) alternatives:

```
@media (prefers-reduced-motion: reduce) {  
  
  *, *::before, *::after {  
  
    animation-duration: 0.01ms !important;  
  
    transition-duration: 0.01ms !important;  
  
  }  
  
}
```



## Highlight errors, don't just describe them

A validation error message at the top of a form that says “Please check the highlighted fields” is useless if no fields are actually highlighted. And even a message like “Email is invalid” is less effective when the user has to scan 15 form fields to find the email input.

Effective validation feedback should include:



- **Field-level highlighting** - the offending input should have a distinct visual treatment: a red border, a tinted background, or both. The user should be able to identify the problem field without reading a single word, purely from the visual change.
- **Animated border or glow** - a subtle pulse or glow animation on the error field draws the eye more effectively than a static color change. CSS `box-shadow` with a red tint and a brief animation (0.3s) works well. This is especially important when multiple fields have errors - the animation helps the user spot them while scanning.
- **Inline error indicators** - place the error message directly below the field it refers to, not in a banner at the top of the page. Inline placement eliminates the need for the user to map “error in email field” to the actual email input.

A simple error glow in CSS:

```
input[aria-invalid="true"] {  
  
    border-color: #ef4444;  
  
    box-shadow: 0 0 0 3px rgba(239, 68, 68, 0.3);  
  
    animation: error-pulse 0.3s ease-out;  
  
}  
  
@keyframes error-pulse {  
  
    0% { box-shadow: 0 0 0 0 rgba(239, 68, 68, 0.5); }  
  
    100% { box-shadow: 0 0 0 3px rgba(239, 68, 68, 0.3); }  
  
}
```

The principle: **message + visual target** always beats message alone. The user should be able to fix the error without any mental mapping or scrolling.



The image shows two side-by-side form panels. The left panel has a red header labeled "BAD" and contains three input fields: "Name" with "John Smith", "Email" with "john@", and "Phone" with "+1 555 123 456". Below the fields is a blue "Submit" button. The right panel has a green header labeled "GOOD" and contains the same three input fields with the same data and a blue "Submit" button.

## Smooth scroll as guidance

When a user submits a form, and there's a validation error on a field that's currently off-screen, what should happen? If the answer is "a red banner appears at the top that says 'Fix errors below'" - the user now has to manually scroll down, find the field, and fix it. That's friction you're imposing on someone who already made a mistake.

Smooth scrolling is useful for:

- **Jumping to validation errors** - after form submission, automatically scroll to the first error field and focus it. The user is taken directly to the problem without having to hunt for it. Combine this with the field highlighting described above for maximum effect.
- **Navigating long forms or pages** - anchor links, table-of-contents navigation, and "Back to top" buttons should all use smooth scrolling. The animation gives the user spatial context - they can see the page moving and understand where they've landed relative to where they were.
- **Revealing results** - after a search or filter action, scroll to the results section so the user doesn't have to wonder where the output appeared.

A typical implementation after form validation:

```
const firstError = document.querySelector('[aria-invalid="true"]');
firstError?.scrollIntoView({ behavior: "smooth", block: "center" });
firstError?.focus();
```



Use smooth scrolling with care - don't hijack the user's scroll behavior or override native scrolling mechanics. Reserve it for intentional, app-initiated navigation (like error focus or anchor jumps), not as a replacement for the user's own scrolling.

The image shows two side-by-side form panels. The left panel, titled 'BAD: BANNER ONLY, NO SCROLL', has a red header. It contains three input fields: 'Full name' with 'Jan Kowalski', 'Email' with 'jan@example.com', and 'Phone' with '+48 600 100 200'. The right panel, titled 'GOOD: SCROLL & FOCUS ON ERROR', has a green header and contains the same three input fields with the same text. The 'GOOD' panel is visually distinguished by its green header and a light green background, suggesting a focus on the error state.

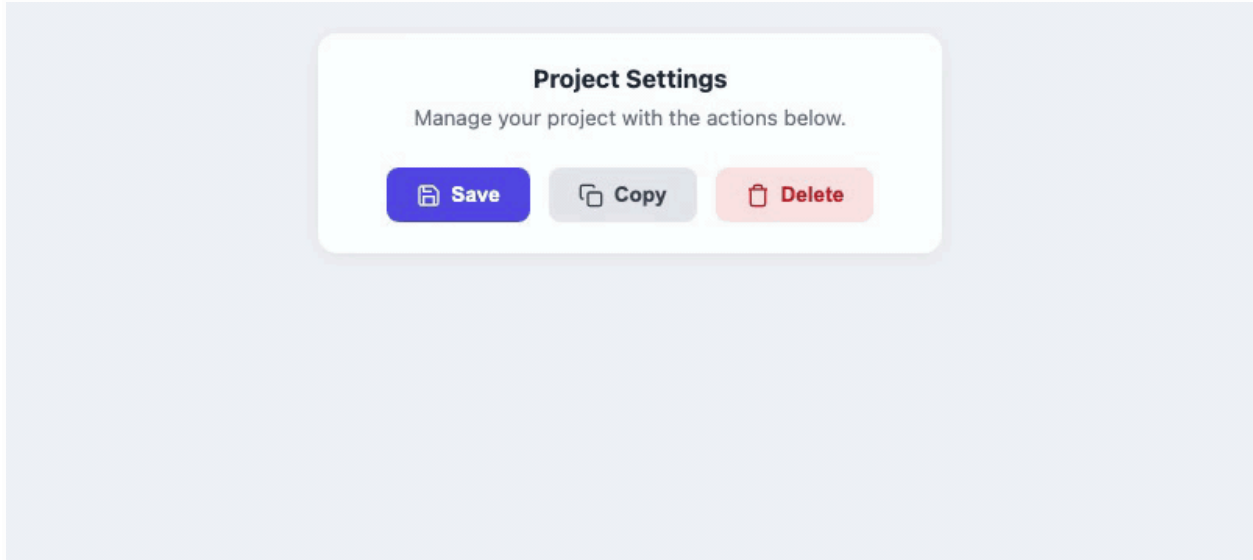
## Toasts as lightweight feedback

After a user performs an action, they need confirmation that it worked. But not every action deserves a full modal or page reload. That's where toasts come in - small, non-blocking notifications that slide into view, deliver their message, and disappear after a few seconds.

Toasts are perfect for:

- **Saves and updates** - "Changes saved" or "Profile updated" as a brief confirmation that disappears on its own. The user doesn't have to click "OK" or close anything.
- **Clipboard copies** - when the user clicks a "Copy" button, a toast saying "Copied to clipboard" confirms the action worked. Without it, the user has to paste somewhere to verify - unnecessary friction.
- **Deletions with undo** - as discussed earlier, "Item deleted - Undo" gives the user a recovery window without interrupting their flow.
- **Background completions** - "Export ready" or "Upload complete" for operations that happened asynchronously while the user was doing something else.





Good toast design follows a few rules: they appear in a consistent position (typically bottom-right or top-right, mirrored in RTL layouts), they stack if multiple appear at once, and they include a close button for users who want to dismiss them immediately.

For informational toasts (“Changes saved”), auto-dismiss after 3–5 seconds is fine. For actionable toasts with an Undo button, use a longer window - 5–10 seconds - to give the user enough time to notice, read, and act. They confirm actions without blocking interaction - the user never has to stop what they’re doing to acknowledge a toast. For screen reader accessibility, toasts need `aria-live` regions - see the Accessibility section for details.

---

## Fuzzy search instead of exact match

When a user types “drk md” into a settings search, they clearly mean “Dark mode.” When they type “notif,” they want “Notification preferences.” But a traditional exact-substring search returns nothing for “drk md” - it only matches if the characters appear consecutively in the original string.<sup>f</sup>

This is a common friction point in search inputs, command palettes, settings panels, and any filterable list. Users don’t think in exact substrings - they think in approximations, abbreviations, and partial recall.

Fuzzy search solves this by matching characters in order, but not necessarily adjacent:

- “**drk md**” matches “**Dark mode**” - all characters appear in sequence, just with gaps.
- “**notif**” matches both “**Notification preferences**” and “**Email notifications**” - ranked by how tight the match is.



- “2fa” matches “Two-factor authentication” - through keyword aliases, even when the literal characters don’t appear in the name.

Implementation tips:

- **Highlight matched characters** - show the user *why* a result matched by highlighting the individual characters that correspond to their query. This builds trust in the search and helps the user refine their input.
- **Rank by match quality** - a match at the start of the string (“Not...” for “not”) should rank higher than a match scattered across the middle. Tighter character clusters score better than spread-out matches.
- **Include keyword aliases** - map common synonyms and abbreviations to items: “2fa” → “Two-factor authentication,” “theme” → “Dark mode.” This catches cases where the user’s mental model uses different words from the label.
- **Use established libraries** - don’t roll your own fuzzy matcher for production. Libraries like [Fuse.js](#) or [uFuzzy](#) handle scoring, ranking, and edge cases out of the box.

The image shows two side-by-side search result panels. The left panel, titled 'BAD: EXACT MATCH ONLY', shows a search bar with '2fa' entered. The results list various settings, but 'Two-factor authentication' is not highlighted. The right panel, titled 'GOOD: FUZZY / TOLERANT SEARCH', shows the same search bar and results, but 'Two-factor authentication' is highlighted with a green background, indicating a successful fuzzy match.

Search Result	Category
Notification preferences	Account
Email notifications	Account
Dark mode	Appearance
Language & region	General
Privacy settings	Security
Two-factor authentication	Security
Keyboard shortcuts	General
Connected accounts	Account



## Avoid modal overuse

Modals are the nuclear option of UI communication. They halt everything - the user can't interact with anything behind the overlay, they must read and respond before continuing. Sometimes that's exactly right (confirming a destructive action on a financial transaction, acknowledging a critical error). But too often, modals are used as the default pattern for every interaction.

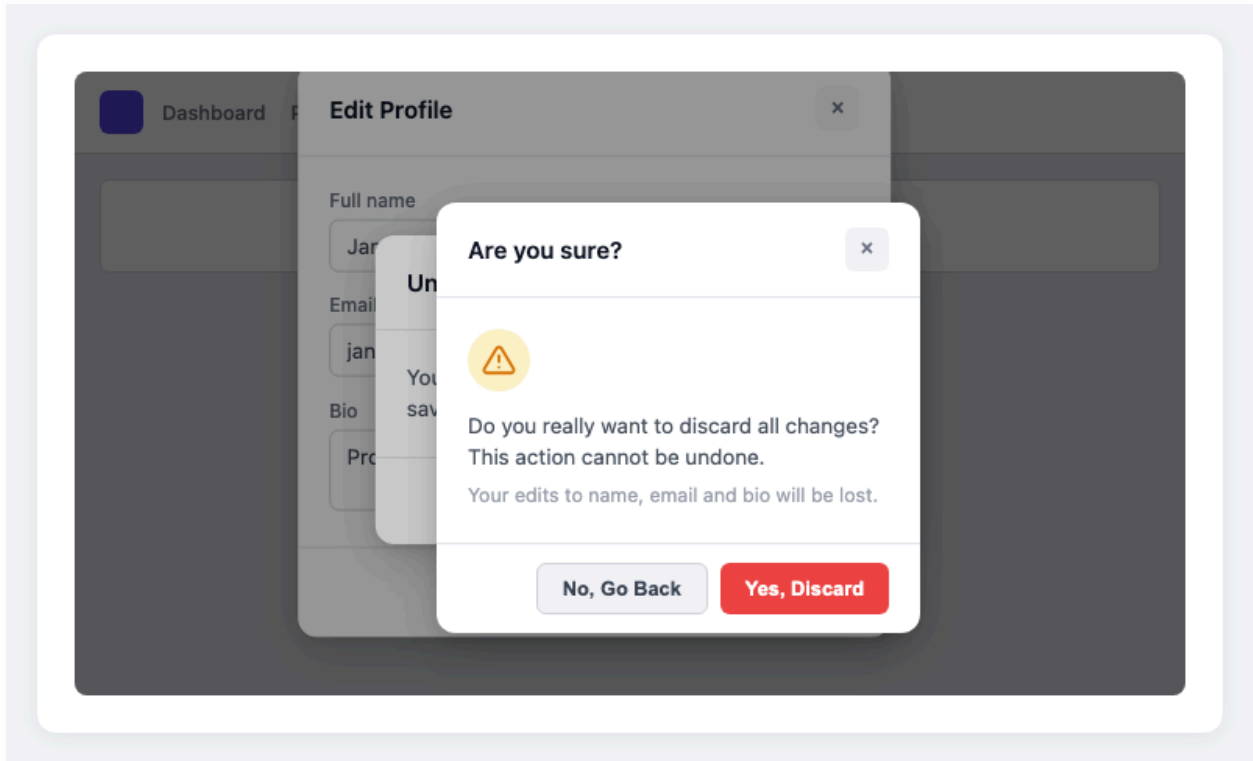
Too many modals:

- **Break context** - every modal is a context switch. The user was focused on a task, but now they're dealing with a dialog. When dismissed, they need to reorient and remember what they were doing.
- **Increase cognitive load** - "Are you sure?" modals force a decision. Confirmation dialogs inside other dialogs (modal inception) create genuine confusion. Each modal is a mental interruption.
- **Slow down workflows** - in admin panels and content management tools, users often perform repetitive bulk actions. If each action triggers a modal, the workflow grinds to a halt.

Before reaching for a modal, consider the alternatives:

- **Inline expansion** - instead of a modal for editing an item, expand the item in-place with inline editing fields. The user stays in context and can see the surrounding items.
- **Drawers and slide-over panels** - a panel that slides in from the right provides a focused editing space without completely blocking the underlying page. The user can still see the list they were browsing.
- **Progressive reveal** - instead of a modal asking for details upfront, reveal additional fields or options gradually as the user interacts with the interface.
- **Undo toasts** - as covered earlier, a "deleted - Undo" toast is almost always better than a "Are you sure?" modal for deletion flows.





## Mobile-first interaction patterns

All the interactions discussed so far need to work across devices - and on mobile, they face a different set of constraints. More than half of all web traffic comes from mobile devices, yet many developer-built interfaces are designed cursor-first and touch-optimized as an afterthought. Mobile isn't a smaller desktop - it has its own interaction model, constraints, and opportunities.

### Thumb zone and reachability

On a phone held in one hand, the user's thumb can comfortably reach only a limited arc of the screen - roughly the lower-center area. The top corners and far edges require stretching or a second hand. This has direct consequences for where you place primary actions.

- **Place primary actions in the bottom half** - navigation bars, FABs (floating action buttons), and key CTAs belong in the thumb's natural arc. iOS and Android both moved their primary navigation to the bottom of the screen years ago for exactly this reason.



- **Avoid top-corner actions for frequent tasks** - a hamburger menu in the top-left corner requires the most uncomfortable reach on a phone. If it's the primary way to navigate your app, consider a bottom tab bar or a bottom sheet instead.
- **Size touch targets generously** - as mentioned in Fitts's Law, 44x44px (Apple) or 48x48px (Google) is the minimum. On mobile, this isn't a suggestion - it's the difference between usable and frustrating.

## Swipe gestures and bottom sheets


Touch interfaces support gestures that have no cursor equivalent. Swipe-to-delete, pull-to-refresh, and swipe between tabs are patterns users already know from native apps - and they expect them on the mobile web too.

- **Swipe actions on list items** - swipe left to delete, swipe right to archive. These reduce the number of taps needed and feel natural on touch devices. Always provide a non-swipe alternative (a button or menu) for discoverability.
- **Bottom sheets instead of modals** - on mobile, a modal centered on the screen wastes the top half (it's unreachable) and forces the user to reach up to close it. A bottom sheet slides up from the bottom, keeping all interactive elements in the thumb zone.
- **Pull-to-refresh** - for content feeds and lists, pulling down to refresh is an expected gesture. If your app shows live or frequently-updated data, implement it rather than relying on a reload button.



**BAD: ACTIONS IN HARD-TO-REACH ZONES**


**NAVIGATION & FAB AT THE TOP**



Key actions require stretching to top of screen

**GOOD: ACTIONS IN THE THUMB ZONE**


**BOTTOM NAV & FAB AT THE BOTTOM**



All primary actions within natural thumb arc

**BAD: CENTERED MODAL ON MOBILE**


**ACTIONS FAR FROM THUMB**



Buttons sit in the hard-to-reach zone

**GOOD: BOTTOM SHEET ON MOBILE**

**ACTIONS IN THUMB ZONE**



All interactive elements within easy reach



# Accessibility is part of UX

Mobile is one dimension of adapting to different users and contexts. Accessibility broadens this to all users, regardless of ability. A significant group of users experiences the web differently from you.

- People with visual impairments use screen readers that parse the DOM and read content aloud.
- People with motor disabilities navigate entirely with a keyboard or switch device - no mouse, no gestures.
- People with cognitive disabilities benefit from clear structure, consistent navigation, and predictable behavior. Color-blind users can't rely on color alone to distinguish states.

If your interface only works for sighted mouse users, it doesn't work. Accessibility isn't a nice-to-have checkbox - it's a core part of UX quality. And in the EU, it's also a legal requirement - the [European Accessibility Act \(EAA\)](#) mandates that many digital products and services meet accessibility standards starting from June 2025. And much of it comes down to writing proper HTML.

## Semantic HTML instead of div soup

The single biggest accessibility win is also the simplest: use HTML elements for their intended purpose.

A `<div>` with an `onclick` handler looks like a button on screen, but to a screen reader, it's just an anonymous container. It's not focusable by default, it doesn't respond to Enter or Space keypresses, and it doesn't appear in the accessibility tree as an interactive element. The user literally doesn't know it's there.

The fix is not to bolt ARIA onto a div to make it behave like a button. The fix is to use `<button>`:

```
<!-- Bad: screen reader sees an anonymous container -->
```

```
<div class="btn" onclick="submit()">Submit</div>
```

```
<!-- Good: focusable, keyboard-accessible, announced as "Submit, button" -->
```

```
<button type="submit">Submit</button>
```

This applies across the board:



- **Navigation** - use `<nav>` and `<a>` for links, not styled divs with click handlers. Screen readers expose `<nav>` as a landmark that users can jump to directly.
- **Lists** - use `<ul>/<ol>` and `<li>` for lists. Screen readers announce "list, 5 items" - giving users context before they dive in.
- **Forms** - use `<form>`, `<label>`, and `<input>` with proper `for/id` binding. A `<label>` connected to an input via `for` means that clicking the label focuses the input (a larger hit area), and screen readers announce what the field is for.
- **Sections** - use `<main>`, `<header>`, `<footer>`, `<aside>`, `<section>` to create landmark regions. Screen reader users navigate by landmarks, the way sighted users scan by visual layout.
- **Tables** - use `<table>`, `<thead>`, `<th>`, `<tbody>` for tabular data. Screen readers let users navigate cell-by-cell and announce row/column headers. A grid of styled divs loses all of that.

The image shows two side-by-side examples of a contact form, illustrating the difference between poor and good HTML semantics.

**Left Panel: BAD: DIV SOUP, NO SEMANTICS**

- Navigation links (Home, Settings, Profile) are styled as buttons but are wrapped in `<div>` tags.
- The main heading "Contact us" is wrapped in `<div class="title">`.
- The subtitle "Send a message" is wrapped in `<div class="subtitle">`.
- The label "Your email" has no `<label>` tag.
- The label "Message" has no `<label>` tag.
- The "Submit" button is wrapped in `<div>`.

**Right Panel: GOOD: SEMANTIC HTML**

- Navigation links are wrapped in `<nav>` and `<a>` tags. The "Settings" link has `aria-current` attribute.
- The main heading "Contact us" is wrapped in `<h1>`.
- The subtitle "Send a message" is wrapped in `<h2>`.
- The label "Your email" is wrapped in `<label for>`.
- The label "Message" is wrapped in `<label for>`.
- The "Submit" button is wrapped in `<button>`.

## Heading hierarchy

Screen reader users often navigate a page by jumping through headings - the way a sighted user scans by scrolling and reading large text. If your headings skip levels or are used for styling rather than structure, this navigation breaks.



Rules:

- **One <h1> per page** - the page title. It's the first thing a screen reader user hears.
- **Don't skip levels** - go h1 -> h2 -> h3, not h1 -> h3. Skipping creates gaps in the heading outline, confusing navigation.
- **Don't use headings for styling** - if you need text to look like a heading but it's not a structural heading, style a <p> or <span> with CSS. Don't abuse <h3> just because you like its default font size.
- **Every content section should start with a heading** - it acts as a navigation anchor for assistive technology users.

You can audit your heading structure with browser devtools. Chrome's accessibility tree view or extensions like HeadingsMap show you the outline as a screen reader would interpret it.

## ARIA - when native HTML isn't enough

ARIA (Accessible Rich Internet Applications) attributes add semantics to elements when native HTML doesn't cover the interaction pattern. But the first rule of ARIA is: **don't use ARIA if a native HTML element can do the job**. ARIA doesn't add behavior - it only adds labels to the accessibility tree. A <div role="button"> is announced as a button, but still isn't focusable or keyboard-operable without extra JavaScript.

When ARIA is the right tool:

- **aria-label** and **aria-labelledby** - provide accessible names for elements that don't have visible text. An icon-only button needs an `aria-label="Close"` because screen readers can't see the X icon.

```
<button aria-label="Close dialog">
```

```
<svg><!-- X icon --></svg>
```

```
</button>
```

- **aria-live** regions - announce dynamic content changes. When a toast appears, or a form validation message shows up, sighted users see it instantly - but screen reader users only know about content that's announced. `aria-live="polite"` waits for the screen reader to finish its current announcement, `aria-live="assertive"` interrupts immediately (use sparingly).

```
<div aria-live="polite" class="toast-container">
```

```
<!-- Toasts injected here are automatically announced -->
```



</div>

- **aria-expanded** - for toggleable elements like accordions, dropdowns, and collapsible sections. Tells screen reader users whether the associated content is visible.

```
<button aria-expanded="false" aria-controls="menu-panel">Menu</button>
```

```
<div id="menu-panel" hidden>...</div>
```

- **aria-invalid** and **aria-describedby** - for form validation. Mark invalid fields and point to their error message so screen readers read both the field label and the error.

```
<input aria-invalid="true" aria-describedby="email-error" />
```

```
<span id="email-error" role="alert">Please enter a valid email</span>
```

- **role="alert"** - immediately announces content to screen readers. Use for error messages and critical notifications that the user must hear right away.

## Hiding elements accessibly

Not all hiding is equal. The CSS property you choose determines whether an element is hidden only visually, or also from assistive technology:

Method	Visible	In layout flow	Announced by screen readers
<code>display: none</code>	No	No	No
<code>visibility: hidden</code>	No	Yes (reserves space)	No
<code>opacity: 0</code>	No	Yes	Yes
<code>.visually-hidden</code> (below)	No	No	Yes



Both `display: none` and `visibility: hidden` hide content from screen readers entirely. This is correct for a UI that is truly inactive (collapsed menus, unopened modals). But sometimes you need content that is invisible on screen yet still available to assistive technology - for example, a descriptive label for an icon button, or a “skip to content” link that only screen reader and keyboard users need.

The standard pattern for this is `.visually-hidden` (sometimes called `.sr-only`):

```
.visually-hidden {  
  
  clip: rect(0 0 0 0);  
  
  clip-path: inset(50%);  
  
  height: 1px;  
  
  width: 1px;  
  
  overflow: hidden;  
  
  position: absolute;  
  
  white-space: nowrap;  
  
}
```

This clips the element to a 1px area offscreen - invisible to sighted users, but still present in the accessibility tree and announced by screen readers. Use it for:

- **Skip navigation links** - “Skip to main content” links that appear on focus (add `:focus` styles to make them visible when tabbed to).
- **Icon-only button labels** - `<button><svg>...</svg><span class="visually-hidden">Close</span></button>` as an alternative to `aria-label`.
- **Form instructions** - additional context for screen reader users that would be visually redundant.

## Keyboard navigation

Every interactive element must be operable with only a keyboard. This means:

- **All clickable elements must be focusable** - native `<button>`, `<a>`, and `<input>` are focusable by default. If you use a `<div>` or `<span>` as a trigger (don’t), you need `tabindex="0"` to add it to the tab order and keyboard event handlers for Enter and Space.



- **Tab order must follow visual order** - users expect Tab to move left-to-right, top-to-bottom. If your CSS layout (flex `order`, grid placement, absolute positioning) visually reorders elements, the tab order still follows the DOM order. Keep DOM order aligned with visual order.
- **Trap focus in modals** - when a modal opens, pressing Tab should cycle through only the modal's interactive elements, not escape to the page behind the overlay. When the modal closes, return focus to the element that opened it.
- **Provide visible focus indicators** - as mentioned in the component states section, never remove `:focus` styles without providing a visible alternative. `:focus-visible` is the modern way to show focus only for keyboard users (not for mouse clicks).

```
/* Focus ring only for keyboard users */
```

```
button:focus-visible {
  outline: 2px solid #4f46e5;
  outline-offset: 2px;
}
```

## Practical checklist

Some quick wins you can apply right now:

- **Add `alt` text to all meaningful images.** Decorative images get `alt=""` (empty, not missing) so screen readers skip them.
- **Don't rely on color alone** - pair color coding with icons, text, or patterns. A form field with a red border also needs an error icon or message.
- **Test with keyboard only** - put your mouse in a drawer for 5 minutes and Tab through your app. You'll find problems in seconds.
- **Check contrast ratios** - WCAG requires 4.5:1 for normal text and 3:1 for large text. Chrome DevTools shows contrast ratios in the color picker. Fix any that fail.
- **Use the `lang` attribute on `<html>`** - screen readers use it to select the correct pronunciation engine. `<html lang="en">` is one line and makes a real difference.
- **Test with a screen reader** - macOS has VoiceOver built in (Cmd+F5). Spend 10 minutes navigating your own app with it. It's the fastest way to find what's broken.

---

## Closing thoughts



Modern tools dramatically accelerate UI creation, but interface quality still depends on implementation details: how progress is shown, how actions respond, how motion behaves, and how stable layouts remain during change.

These small technical decisions compound into a noticeably better user experience - and they're fully within reach at the frontend implementation level. You don't need a design degree or a UX certification.

You need to care about the details that users feel but rarely articulate, and to treat every loading state, every transition, and every error as an opportunity to make your interface feel like it was crafted with intention.

*Special thanks to Jakub Łyko, Łukasz Pyrzyński, Marcin Baraniecki, Dariusz Jaroń for their feedback and peer review of this blog post.*

