Designing
Loading and Progress States

- Loading states are critical UX communication tools — not decoration
- Uncertainty and lack of feedback increase perceived wait time and frustration
- Core goal: set clear expectations, provide meaningful feedback, reduce perceived time
- Every extra second of load time reduces conversions by ~7%; 63% bounce rate at 4+ seconds
- Course roadmap: taxonomy, patterns, accessibility, AI-era design, measurement

The Cost of Waiting:
User Behavior and Business Impact

- 1–2s load: 9% bounce rate; 3–5s load: 38% bounce; 5s+: 53%+ bounce
- 53% of mobile users abandon sites taking more than 3 seconds to load
- Each extra second of load time (0–5s range) reduces conversions by ~4.4–7%
- BBC loses 10% of audience per added second; Vodafone's 31% LCP gain drove 8% more sales
- Speed is a business metric — loading UX directly shapes revenue, trust, and retention

Taxonomy of Loading States

Determinate vs. indeterminate: measurable progress vs. unknown duration waits

Passive vs. active: system-driven delays vs. user-initiated actions

Scope levels: full page, sectional components, and inline elements

Related patterns: background ops, optimistic UI, streaming, deferred content

Match indicators to expected duration, known progress, and user context

Progress Indicators: Determinate and Indeterminate Patterns

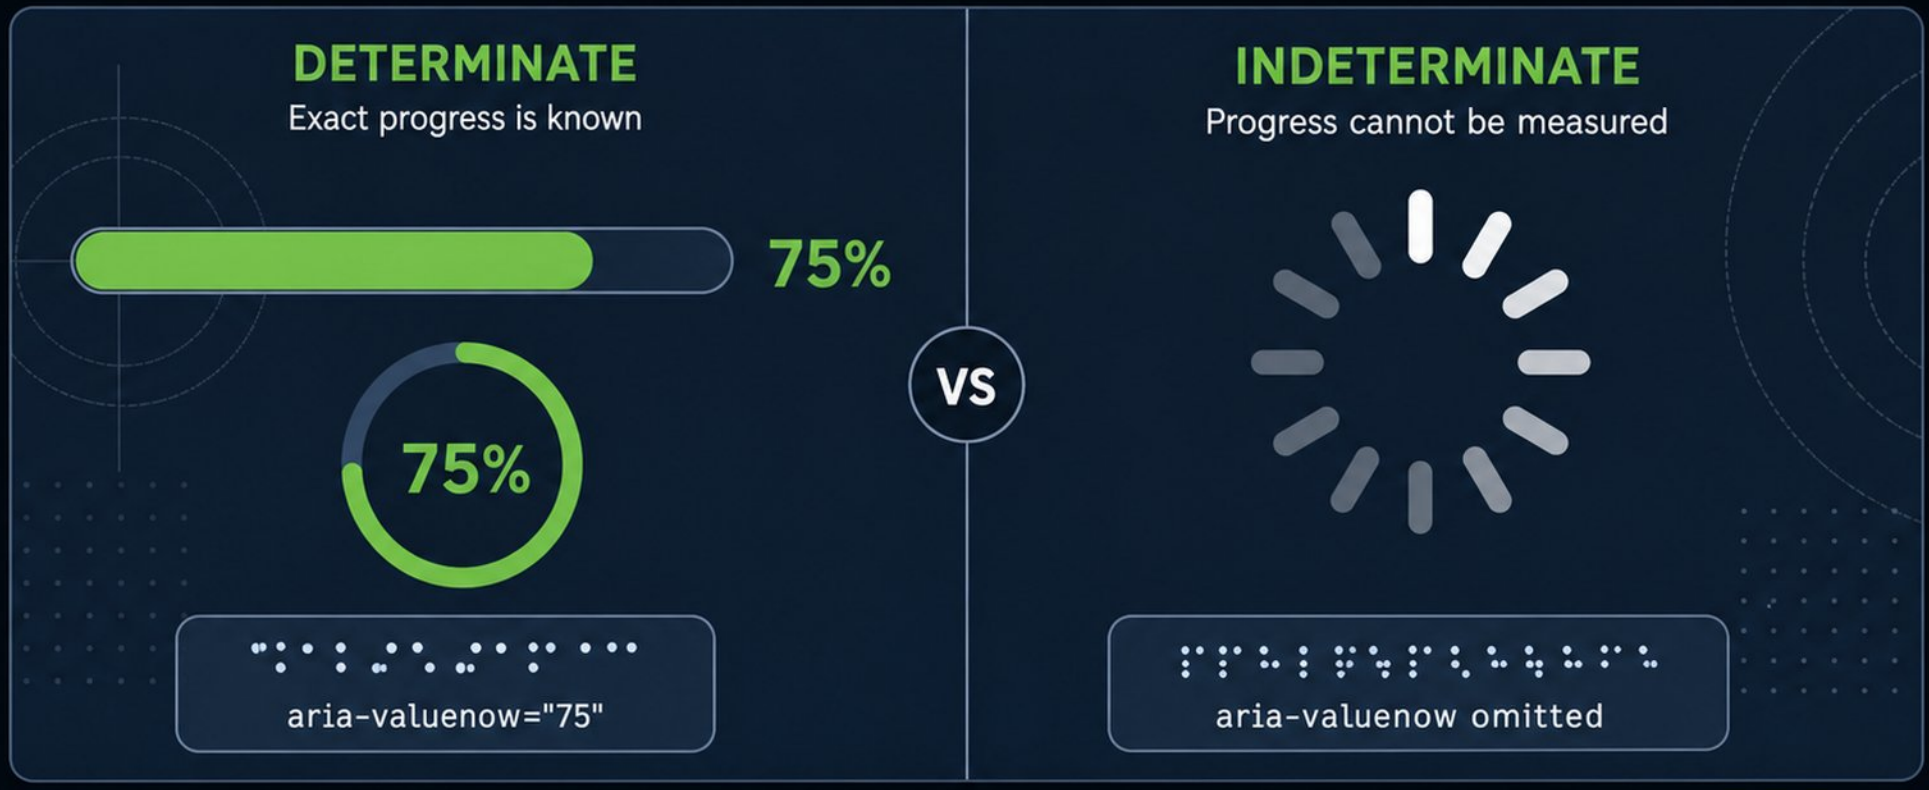

- Determinate: linear/circular bars, percentages, and step counters showing exact progress
- Indeterminate: spinners, shimmers, and looping animations when progress cannot be measured
- Always label tasks with `aria-label` or `aria-labelleddy` — never leave progressbars unlabeled
- ARIA contract: determinate uses `aria-valuenow`; indeterminate omits it entirely
- Use native `<progress>` element as a semantic fallback when custom styling isn't required

Skeleton Screens and Placeholder UI

- Skeleton screens feel ~20% faster than spinners by setting spatial expectations
- Matching the final layout prevents jarring shifts when content loads
- Pulse or wave animations at 300–700 ms cycles signal ongoing progress
- Combine with streaming for smooth skeleton-to-text transitions
- Use role="status" to announce loading start and completion

Empty, Error, and Interruption States

- Empty states: show what belongs here and how to get started
- Error recovery: timeout messages, retry actions, preserve user input
- Interruption handling: network loss, task persistence, stale-while-revalidate
- Tone of voice: clearly distinguish "still working" from "something went wrong"
- Graceful degradation: display succeeded and failed items separately

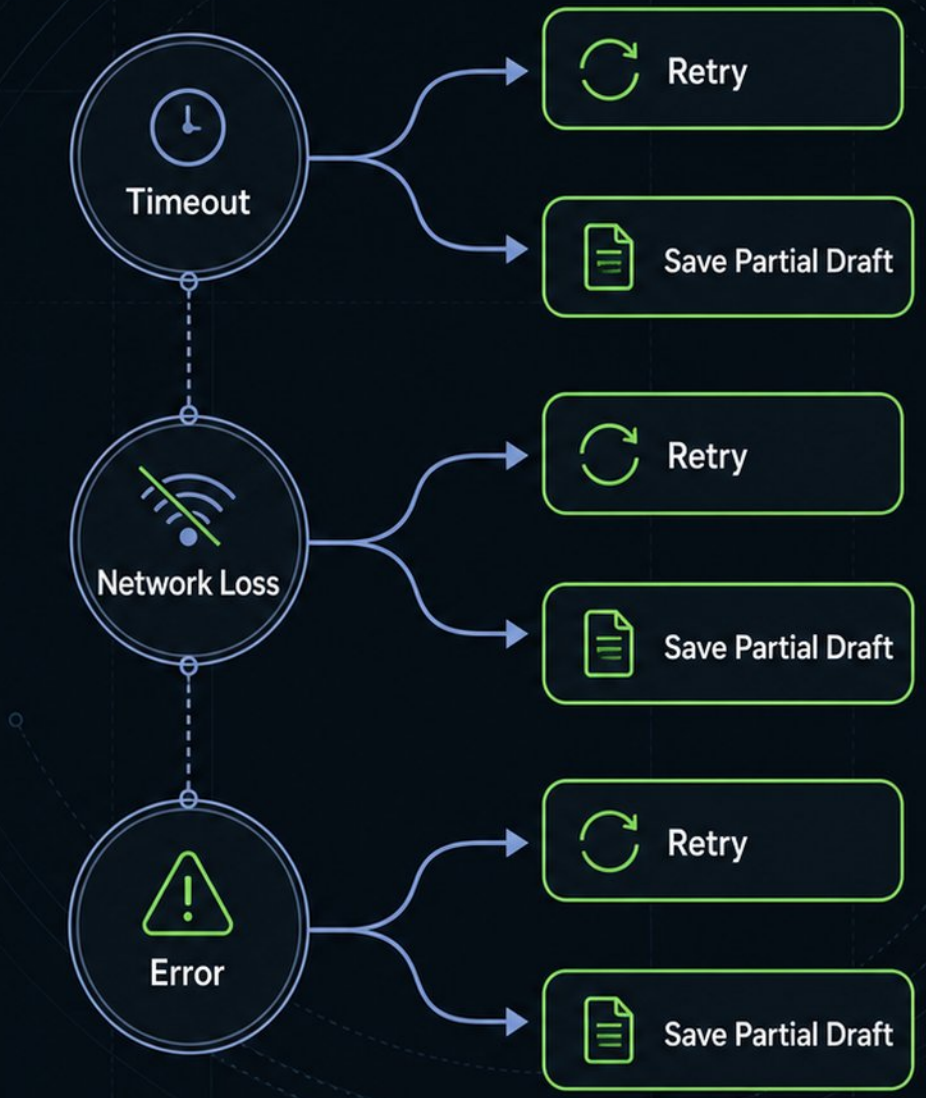

Perceived Performance: Making Waits Feel Shorter

- Actual vs. perceived load time: perception is what users act on
- Feedback within 100ms; animations with start/end acceleration
- Skeleton screens feel ~20% faster than spinners for same wait
- Give users something useful: contextual messages, tips, partial previews
- Optimistic UI: act immediately on intent while system catches up

Accessible Loading Design: ARIA, Focus, and Motion

- Name every progress indicator: use aria-label or aria-labelledby — WCAG 4.1.2 requires it
- Announce state changes via role="progressbar", role="status", and polite live regions
- Keep focus logical during content replacement; never trap users in loading states
- Respect prefers-reduced-motion with static fallbacks for indeterminate animations (WCAG 2.3.3)
- Beyond 1 second add contextual text; beyond 4 seconds offer retry or fallback actions

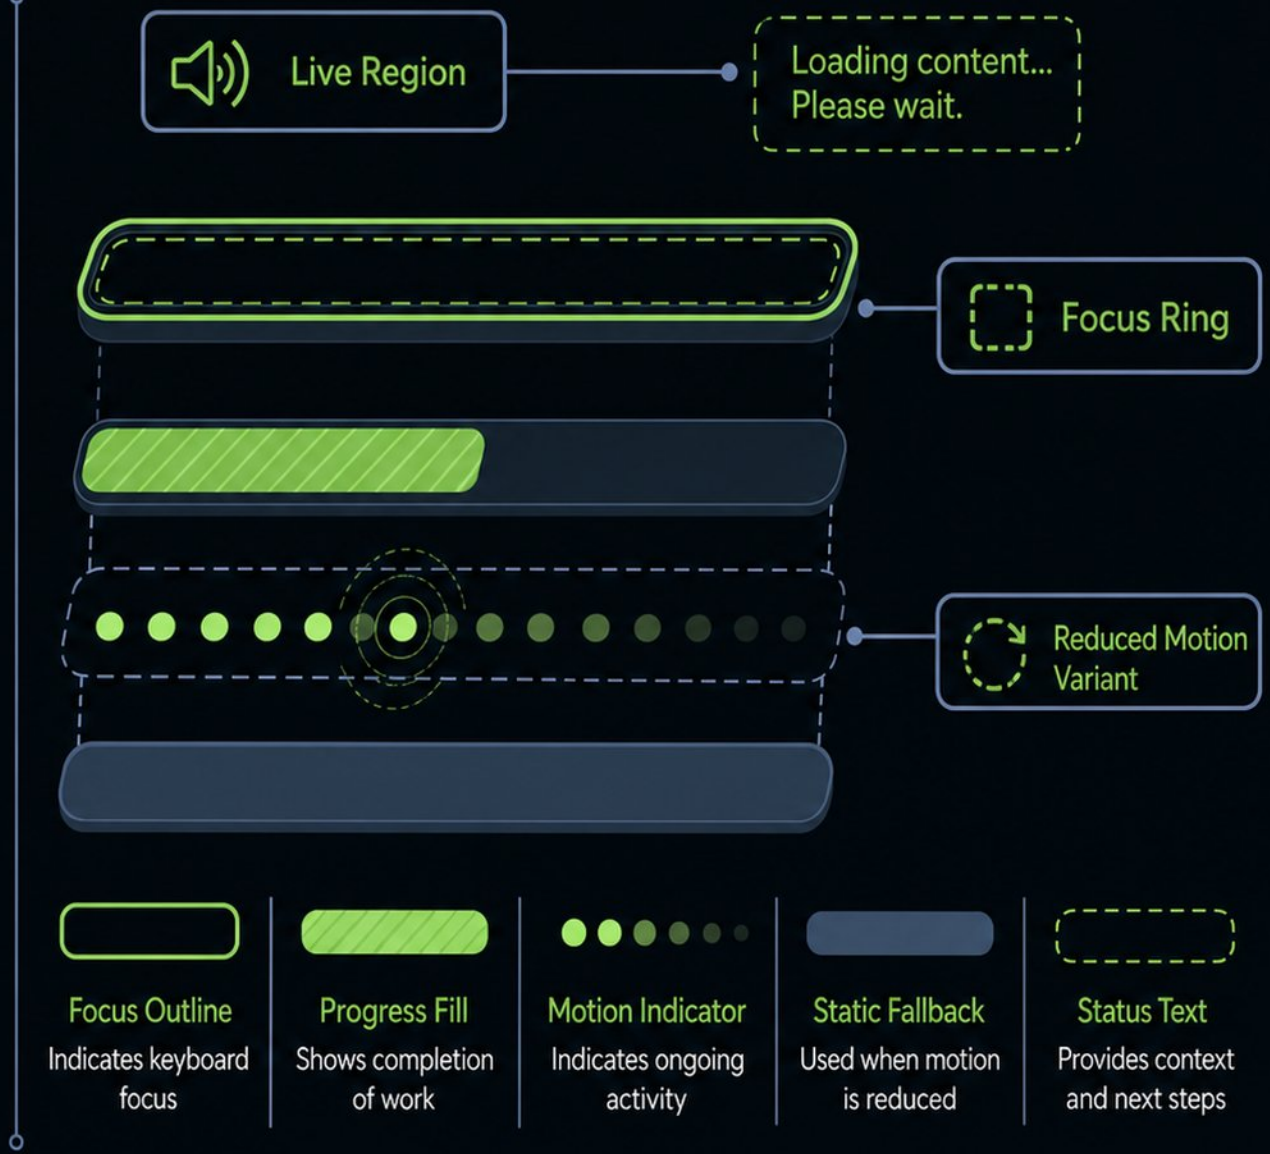

Designing for AI Latency: The New Frontier

- Match feedback to latency zones: instant, short (2-10s), medium (10-60s), and background (60s+)
- Fill the Time to First Token gap with thinking indicators and skeletons
- Stream tokens incrementally with typing cursors and reserved space to prevent layout shifts
- Use plan previews and dynamic checklists to show steps, mark completions, and track real progress
- Provide a working Stop button that cancels generation, saves cost, and gives users control

AI Transparency Patterns: Building Trust Through Visibility

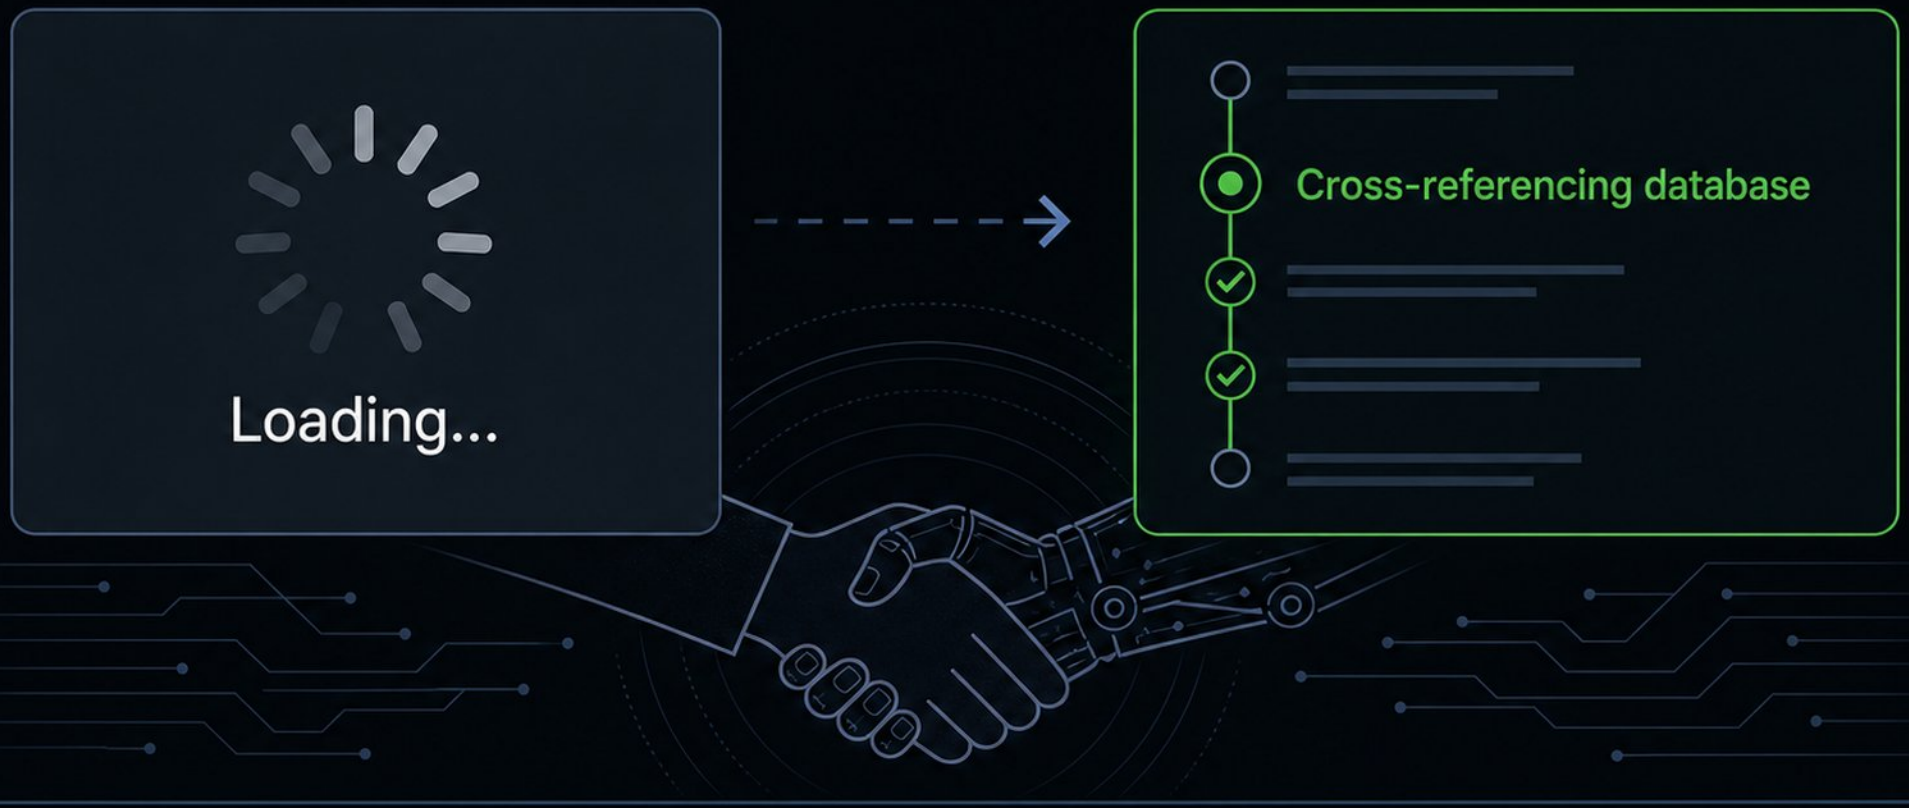

- Replace "Loading..." with precise actions like "Cross-referencing knowledge base"
- **Living Breadcrumb:** subtle status updates for low-stakes background work
- **Dynamic Checklist:** list all steps, highlight current, mark completed phases
- **Thinking Toggle:** collapsed-by-default reasoning traces for verification
- **Audit Trail:** persistent log of system decisions surviving task completion

Real-World Patterns Across Domains

- **E-commerce:** checkout flows and tracking — speed directly shapes cart abandonment and revenue
- **SaaS and dashboards:** data fetching and export generation — progress bars with background notifications
- **Mobile apps:** app launch and pull-to-refresh — skeleton screens and OS-native indicators
- **AI and generative features:** streaming responses and thinking indicators — named tool calls and plan previews
- **Cross-domain rule:** size the indicator to the wait, name the work, and always offer cancel or background

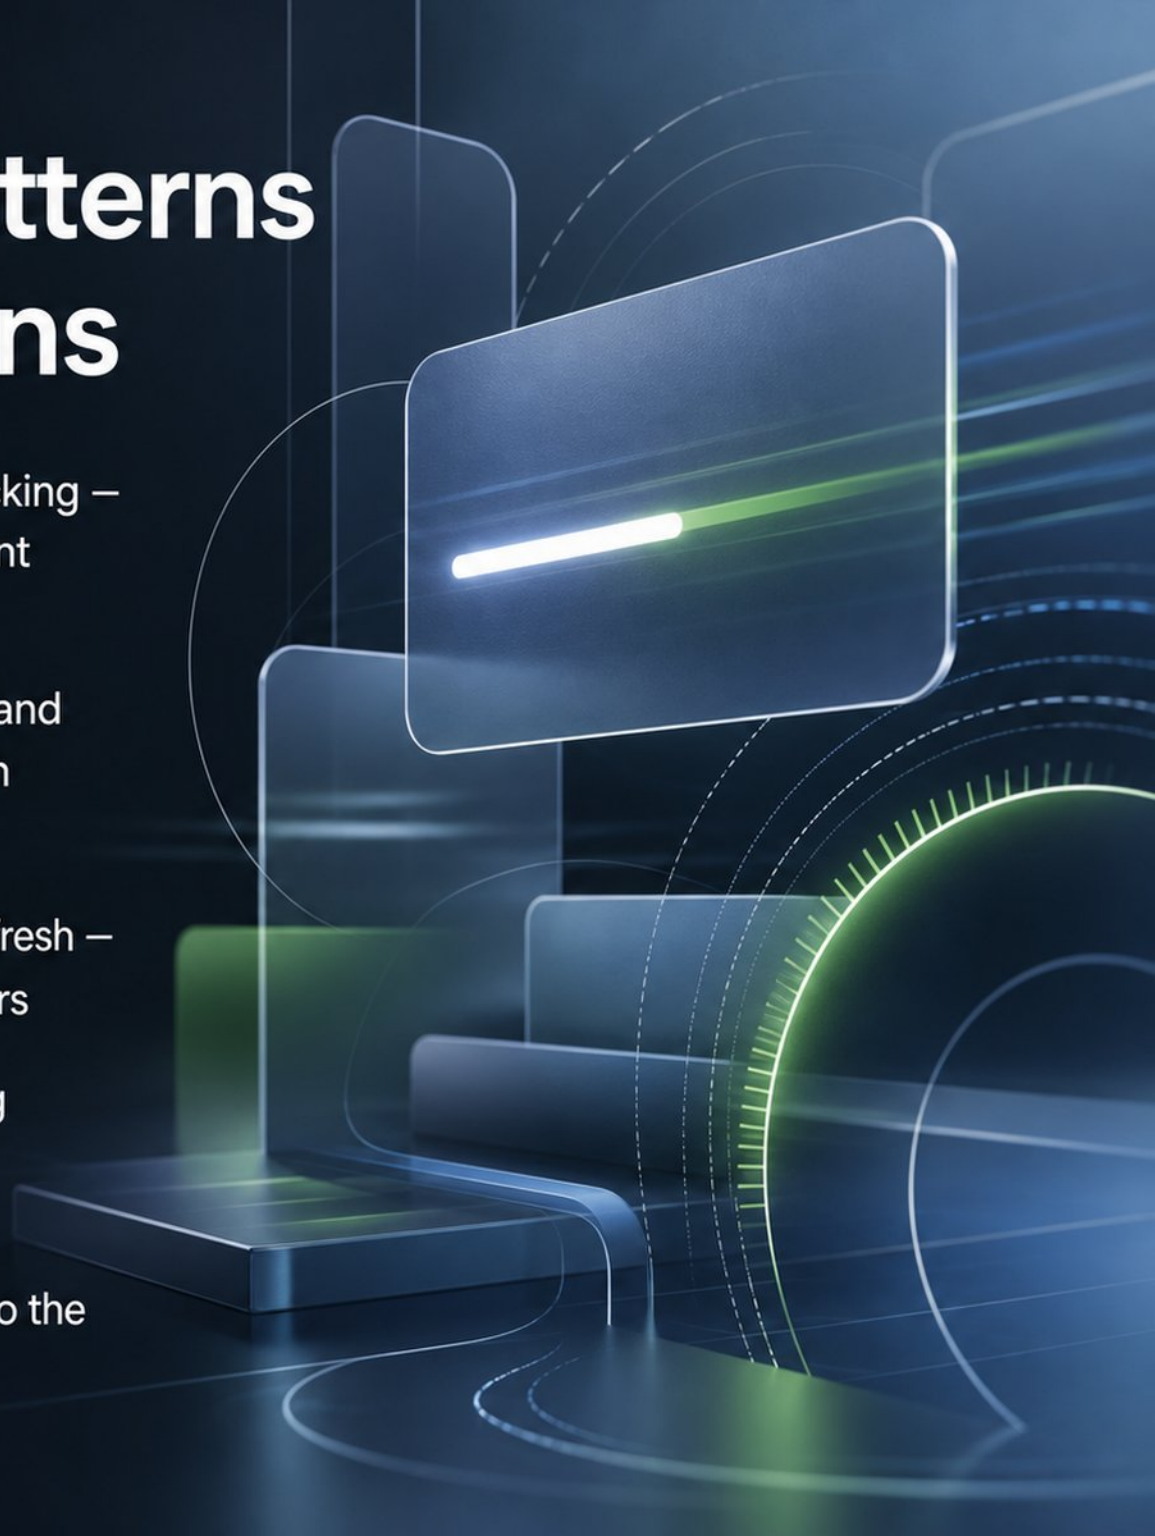

Measuring and Testing Loading Experiences

- Quantitative metrics: TTI, LCP (<2.5s), Speed Index, and Core Web Vitals thresholds
- Qualitative methods: user sentiment surveys, task completion rates, and perceived latency studies
- RUM: calibrate progress indicator accuracy against actual timing data from real sessions
- A/B test loading state variations to measure conversion and abandonment differences without bias
- Testing tools: Lighthouse, WebPageTest, screen reader verification, reduced-motion, and color contrast checks

Design System Integration and Governance

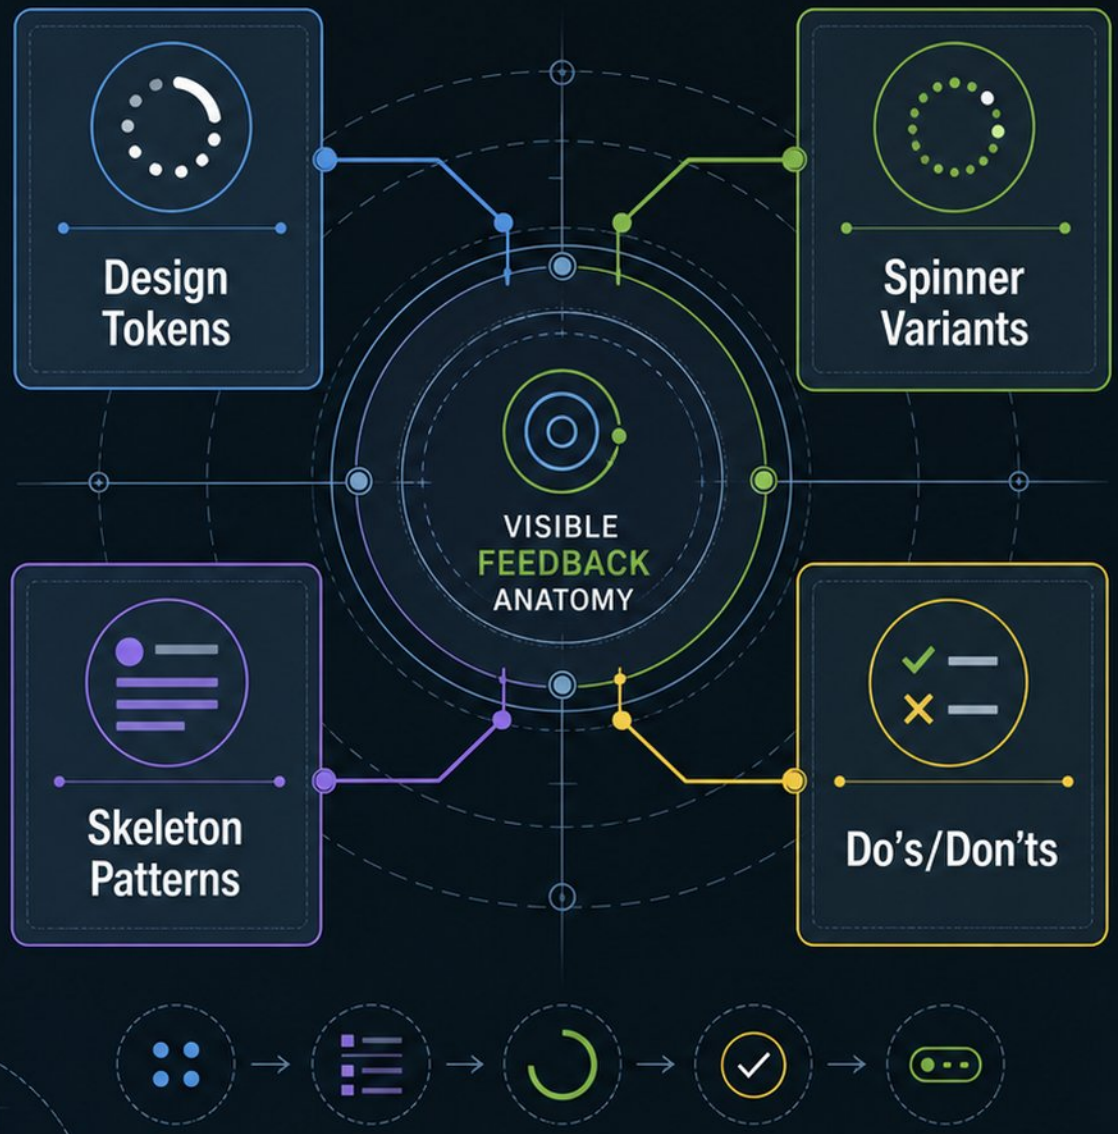

- Define design tokens for animation duration, easing, skeleton shapes, and color contrast
- Build a component library: spinner, skeleton, progress bar, and inline status badge
- Set usage guidelines: pattern selection, min/max display times, and escalation rules
- Establish API contracts for progress data, loading handoff, and edge-case mapping
- Enforce governance with lint rules, review criteria, dos/don'ts, and code playgrounds

PersonWise.

PersonWise • AI digital-human course platform

Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/75a086bc/public