CS-639 — Interaction Design Studio

Design Systems for Intelligent Products*

Professor Bilge Mutlu

* Visuals obtained from the web, used for educational purposes without modification. Attribution on last slide.

Today

— **W12 is reflection-only** — this is the only lecture this week
— **Wed and Fri** are A2 work time — use them well
— **Today:** How to systematize what you are building in A2
— Design systems give you a vocabulary for the components intelligent systems need — and a way to make them consistent, reusable, and legible

You have been designing intelligent interactions for four weeks. Today is about making those designs systematic.

Where We Are

A2 is due in one week. Today is about systematizing what you already have — not starting something new.

What You Have (W08–W11)	What Today Adds
Agency decisions	Components that make those decisions **visible**
Timing decisions	Microinteractions that **communicate** timing
Collaboration patterns	Controls that enable smooth **handoffs**
Context-awareness	Indicators that show what the system **knows**

Design systems turn individual design decisions into a coherent language.

Part 1: What Is a Design System?

Reusable components + rules for how they work together

Design System Basics

A design system is a **shared language** of reusable components and rules (Kholmatova, 2017). Industry convention organizes them into layers:

Layer	What It Contains	Source
Tokens	Named values: colors, spacing, type	Salesforce Lightning (2014)
Components	Reusable UI elements built from tokens	Frost (2016) — Atomic Design
Patterns	How components compose into flows	Kholmatova (2017) — functional + perceptual

A design system is not just a component library — it is a shared language that makes design decisions consistent, scalable, and communicable.

Material Design

Google's Material Design shows what a mature design system looks like:

— Systematic tokens: color, elevation, shape, motion

— Extensive component library with defined states (default, hovered, focused, pressed, disabled)

— Cross-platform consistency

— Every component is documented with usage guidelines, anatomy, and specs

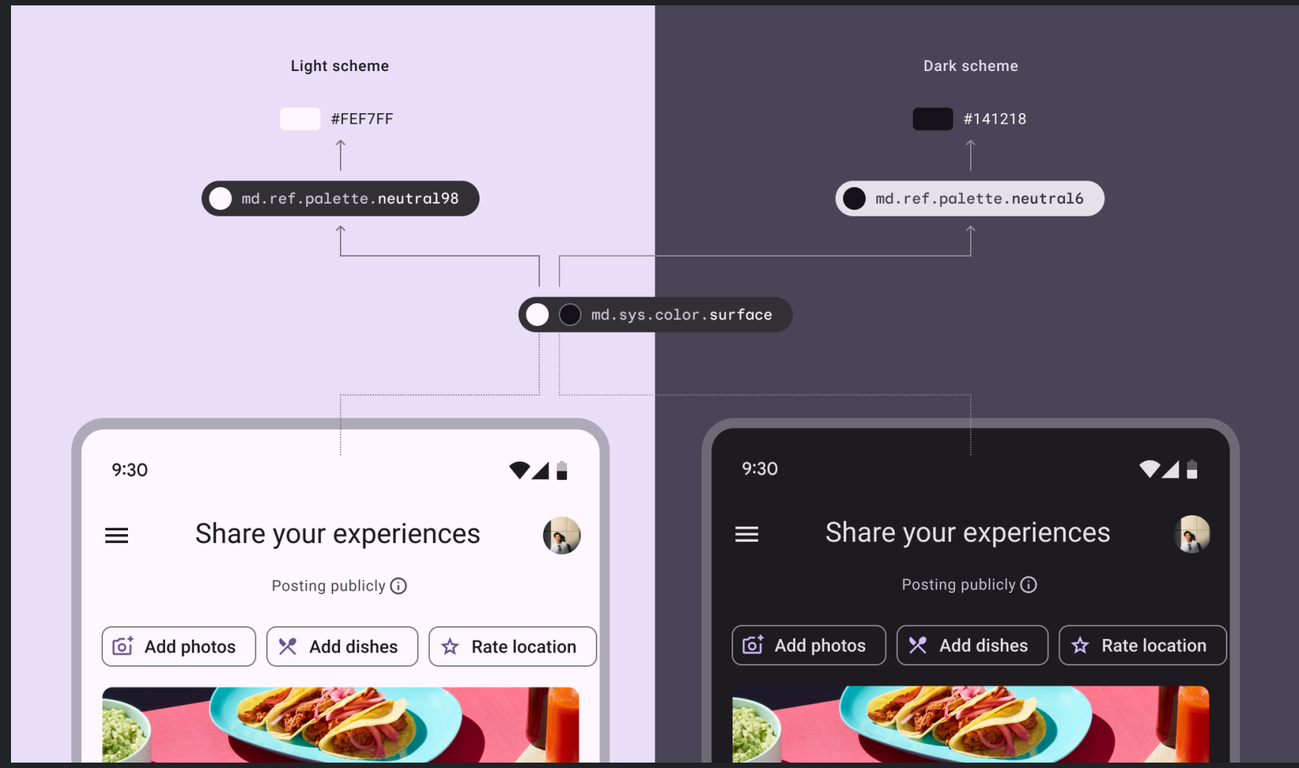

Material Design handles standard UI well. But what does a "confidence indicator" component look like? It is not in the library.

Apple HIG

Apple's HIG takes a platform-specific approach:

— Adaptive layouts that respond to device and context
— Accessibility built into every component
— Strong opinions about when to use which pattern
— Some Siri and Apple Intelligence guidance — but limited

Both Material Design and Apple HIG handle standard UI. Neither fully addresses **intelligent components** — processing states, explainers, override controls, adaptation indicators.

Standard design systems were not built for AI. Intelligent products need new component types.

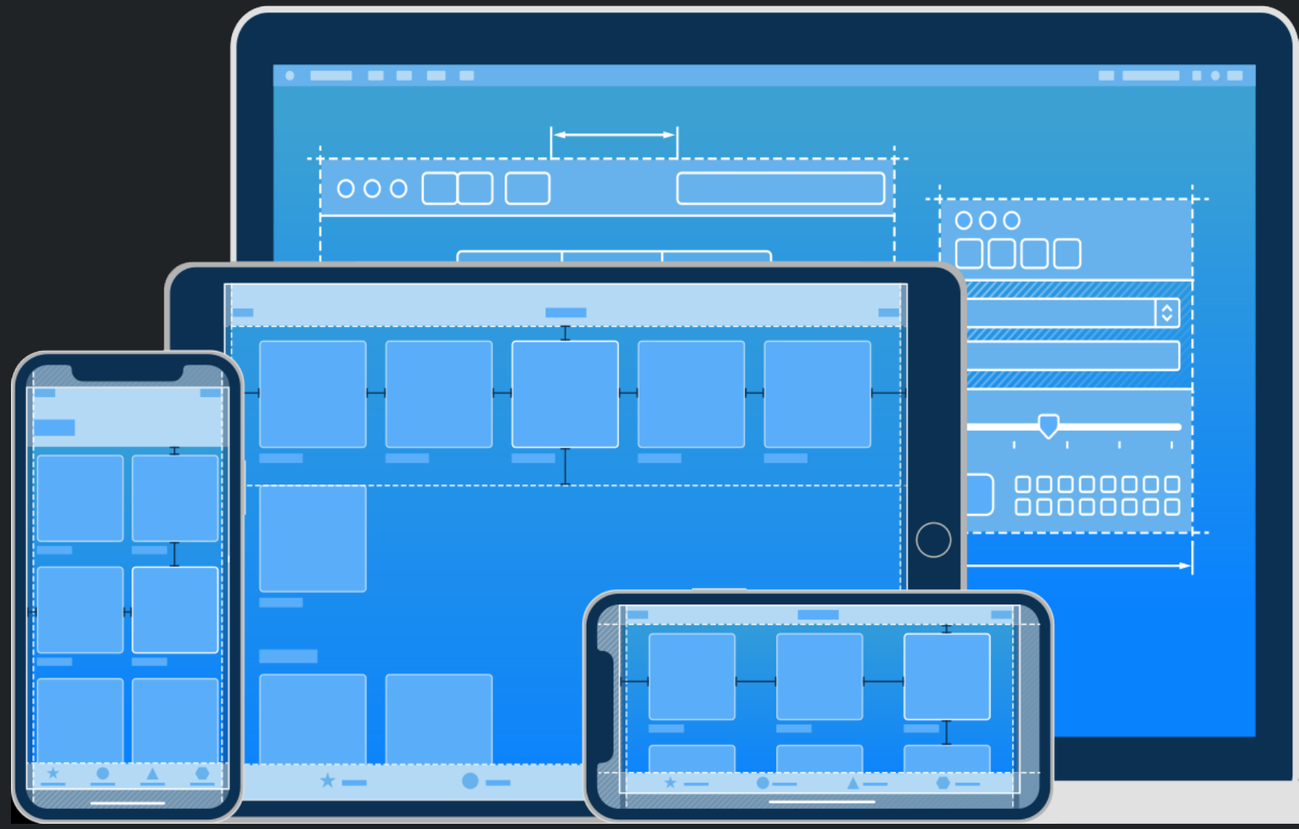

Part 2: Intelligence-Specific Components

Components that standard design systems don't cover

Five Component Types for Intelligent Systems

Operationalizations of Amershi et al.'s (2019) HAX guidelines and Weisz et al.'s (2024) generative AI principles as UI components:

Component Type	What It Does	Principles Served	Material Property
Confidence indicators	Show how sure the system is	P4 (Communicate Uncertainty)	Agency, Context (SA levels)
Processing states	Show the system is "thinking"	P3 (Honest Capabilities), P7 (Explain When Needed)	Proactivity, Collaboration
Explainers	Show why the system did something	P3 (Honest Capabilities), P7 (Explain When Needed)	All four
Override controls	Let users redirect or correct	P9 (User Correctability), P10 (Smooth Transitions)	Agency, Collaboration (OPD)
Adaptation indicators	Show what changed and why	P6 (Contextual Awareness), P8 (Balanced Initiative)	Context-Awareness

These five types are the building blocks of intelligent UI — each one turns a design principle into something the user can see, understand, and act on.

Confidence Indicators

P4 (Communicate Uncertainty) + SA levels (W11)

— Progress bars, probability labels, badges, color-coded levels

— Low-stakes → subtle: "About 1,230,000 results"

— High-stakes → explicit: "85% confidence — benign lesion"

Design rule: Stakes determine how much confidence to show.

Where does your A2 system make predictions? Can the user see how confident the system is?

AI Confidence Colors

Green for high confidence (≥85%)

Yellow for medium confidence (60–84%)

Red for low confidence (<60%)

Abdul-Rashid Lansah Adam
Product Designer

Processing States

P3 (Honest Capabilities) + OPD Observability (W10)

— A loading spinner says "wait." AI processing should say **what** it is doing.
— "Loading..." → "Analyzing your glucose patterns..."
— Examples: ChatGPT's "Searching the web...", Figma AI's "Generating..."

Design rule: The more the user needs to trust the output, the more transparent the processing state should be.

When ChatGPT shows "Searching the web..." before answering, how does that change your trust compared to when it just starts typing?

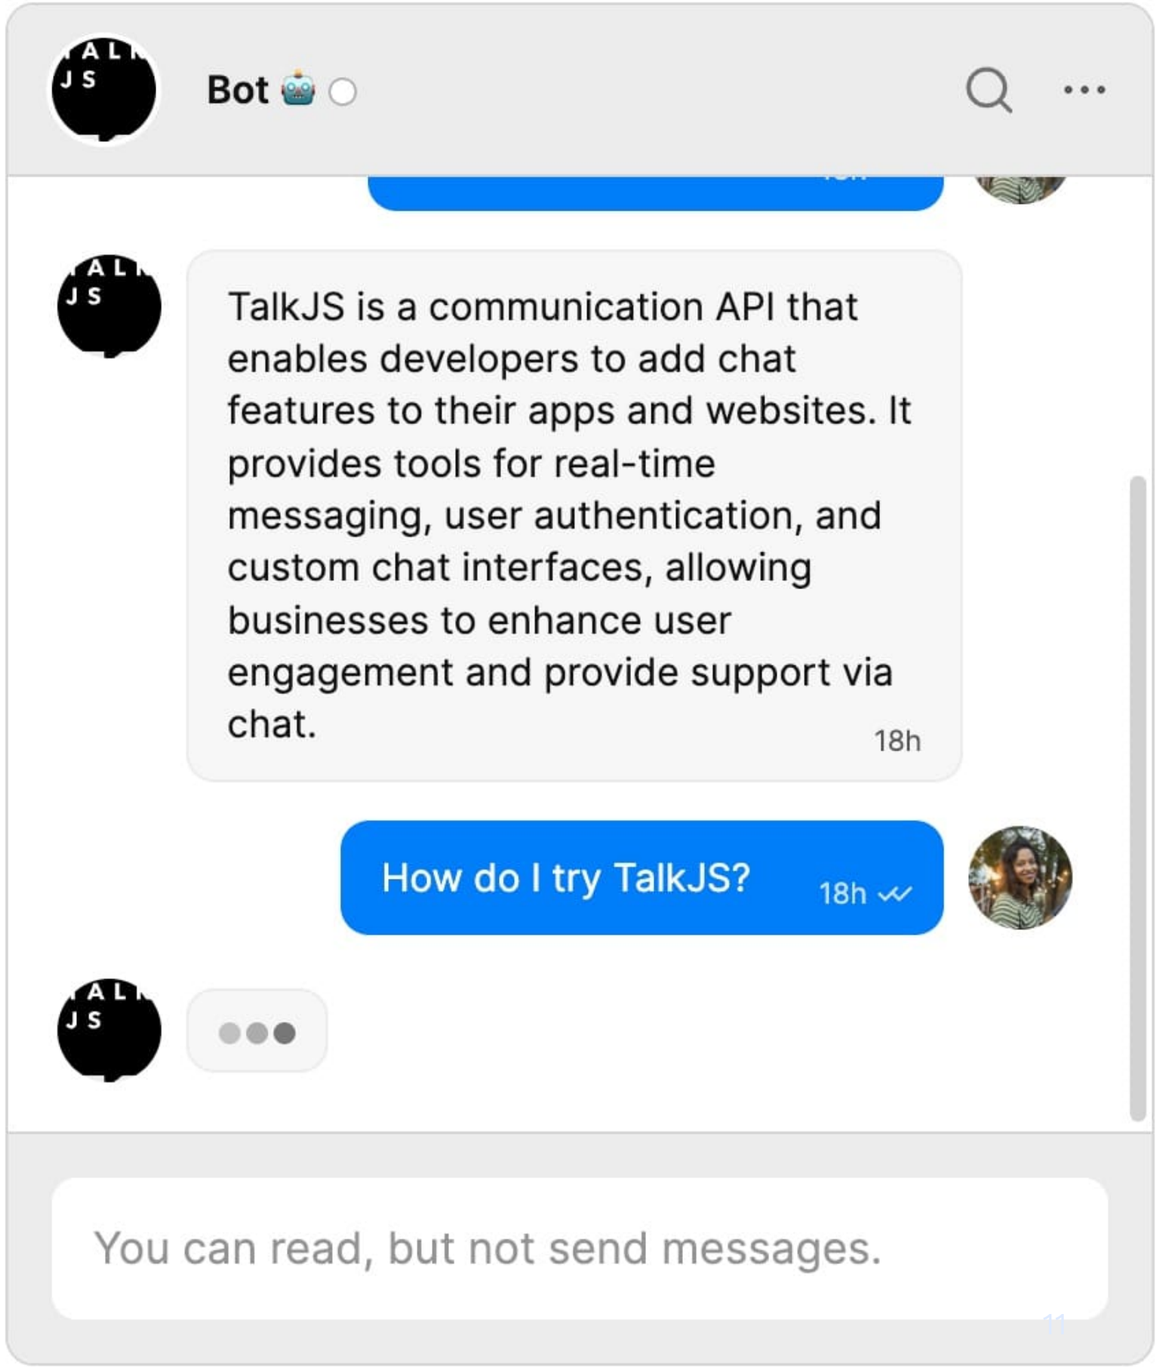

Explainers

P3 (Honest Capabilities) + P7 (Explain When Needed) + SA levels (W11)

Three levels, matched to stakes:

— **Glanceable:** icon or badge — "Made for you"
— **Summary:** one sentence — "Based on your afternoon glucose pattern"
— **Detailed:** full reasoning on demand — for high-stakes, consequential actions

Design rule: Match explanation depth to stakes and user expertise.

Spotify Daylist's label is a glanceable explainer. A medical diagnosis needs a detailed one. Match the depth to the stakes.

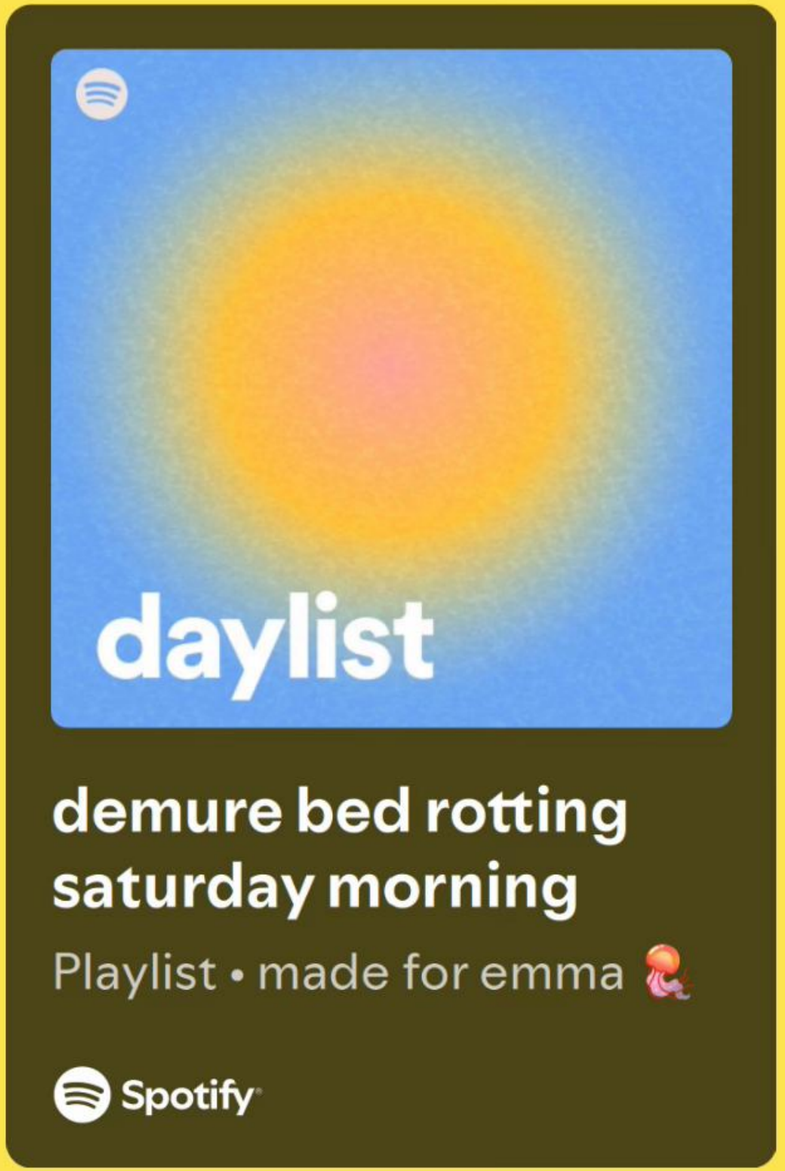

lemon8
@emmabeflippin

Override Controls

P9 (User Correctability) + P10 (Smooth Transitions) + OPD Directability (W10)

— **Dismiss** — "not now" (suggestion returns later)
— **Correct** — "that's wrong" (system learns)
— **Stop** — "don't do this again" (suppress permanently)
— **Undo** — "take it back" (reverse the action)

Design rule: The override must match the action's reversibility. Irreversible actions need confirmation, not just undo.

Gmail "Undo Send" works because sending is briefly reversible. A financial app auto-transferring money needs confirmation before acting.

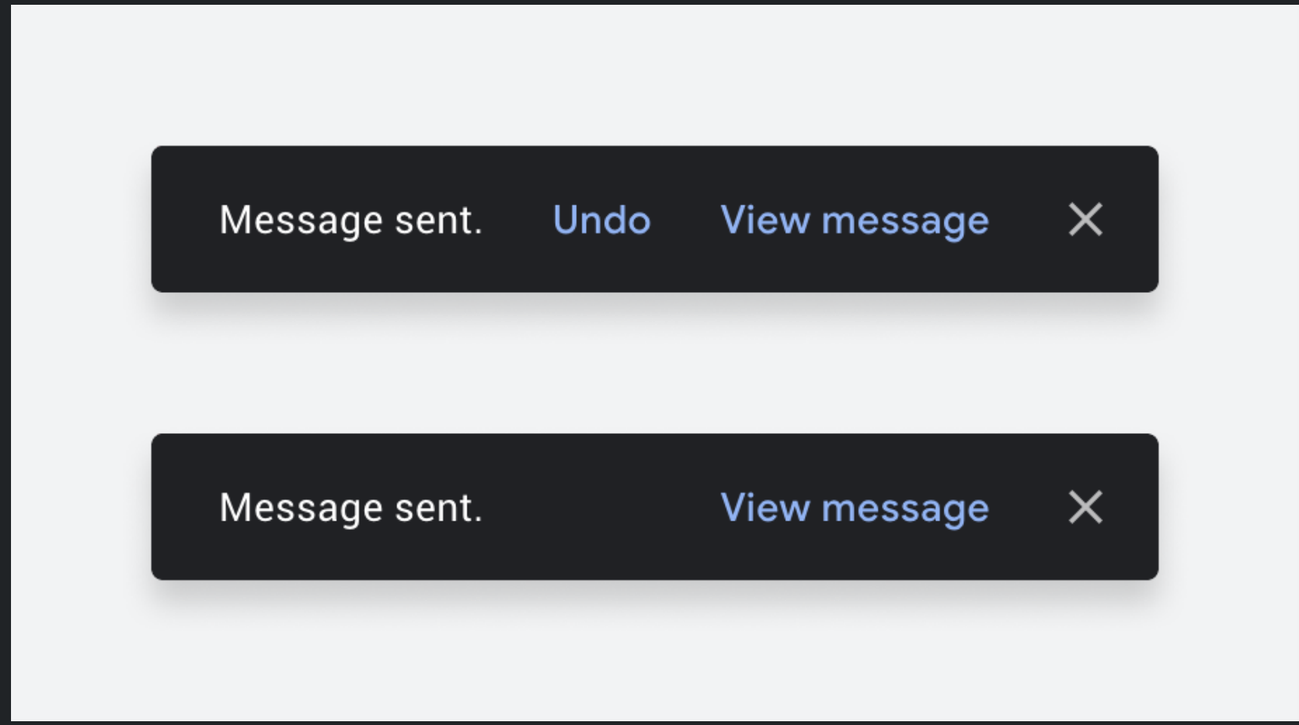

Adaptation Indicators

P6 (Contextual Awareness) + P8 (Balanced Initiative) + Context (W11)

Silent adaptation = confusion. Make it visible:

— **Transition animation** — user sees the change happen
— **Context label** — "Adapted for your morning routine"
— **Before/after** — "Previously: evening recommendations"
— Examples: iOS Focus Mode on lock screen, Spotify Daylist labels

Design rule: If the user can't tell the system adapted, the adaptation is invisible — and potentially confusing.

In your A2 system: when the interface changes based on context, can the user see what changed and why?

Part 3: Microinteractions for Intelligent Feedback

How intelligent components behave moment to moment

Microinteraction Anatomy

Dan Saffer's (2013) framework, adapted for intelligent UI:

Element	Standard UI	Intelligent UI
Trigger	User taps a button	System-initiated (W09 Act/Ask/Wait)
Rules	If valid, submit	Inference, confidence thresholds
Feedback	Success/error	Certainty level + override option
Loops	Repeat or change state	System **learns** from response

In standard UI, the microinteraction ends. In intelligent UI, it feeds back into the system.

AI Microinteraction Patterns

Four patterns that cover most intelligent interactions:

Pattern	Trigger	Components Used
Suggestion	System detects opportunity	Explainer + Override
Confidence warning	Uncertainty exceeds threshold	Confidence indicator
Context shift	Environment/task changes	Adaptation indicator
Collaborative handoff	Leadership shifts (OPD)	Processing state + Override

If you can design these four patterns well, your A2 system will feel coherent.

Timing and Tone

Timing: Too fast = intrusive. Too slow = useless. Appear at a natural pause; exit gracefully when dismissed.

Tone: Match to confidence and stakes:

	Low Stakes	**High Stakes**
High confidence	Direct: "Meeting moved to 3 PM"	Direct + evidence: "Critical low in 20 min"
Low confidence	Tentative: "You might like this"	Tentative + explain: "This could indicate X — consult your provider"

High confidence + high stakes = direct. Low confidence + high stakes = tentative with explanation. Never be confident about something uncertain.

Part 4: Applying to A2

Using today's frameworks to strengthen your assignment

Audit Your A2 System

Quick self-check — pick 2-3 moments where your A2 system acts intelligently:

1. Does the user know how confident the system is? → **Confidence indicator**
2. Can the user see what the system is doing? → **Processing state**
3. Can the user understand why? → **Explainer**
4. Can the user redirect or correct? → **Override control**
5. Can the user tell the system adapted? → **Adaptation indicator**

If any answer is "no" for a high-stakes moment, that is a gap in your design.

Component–Principle Mapping for A2

Annotate your A2 screens with the component types and principles they serve:

In your A2 system...	Component Needed	Principle
System makes a prediction	Confidence indicator	P4 (Communicate Uncertainty)
System processes or generates	Processing state	P3 (Honest Capabilities), OPD Observability
System acts autonomously	Explainer + Override control	P3, P7 (Explain), P9 (User Correctability)
System adapts to context	Adaptation indicator	P6 (Contextual Awareness), P8 (Timely Intervention)
User and system collaborate	Handoff indicator	P10 (Smooth Transitions), OPD
System learns from user	Feedback acknowledgment	P9 (User Correctability), P12 (Productive Imperfection)

You do not need all five component types. Identify which ones your system needs most — and make those excellent.

This Week

A2 refinement and reflection

This Week

— **Today** was the only lecture — Wednesday and Friday are A2 work time
— Use Wednesday and Friday to **refine your A2 screens and flows** using today's frameworks
— **Office hours** available for feedback — bring specific questions about your component choices
— **A2 due Monday, April 20**

Day	What Happens
Monday (today)	Design systems lecture
Wednesday	A2 work time — refine screens and flows
Friday	A2 work time — polish and document

You have one week. Use it to systematize what you have, not to start over.

Reflection: Design Systems Audit

Due before next Monday | Graded: check system

Look at your A2 system and respond to the following:

1. **Component inventory:** Which of the five intelligence component types (confidence indicators, processing states, explainers, override controls, adaptation indicators) does your system need? Why?
2. **Deep dive:** Pick the most critical component type for your system and describe it in detail — what does it look like? What principle does it serve? How does it behave when triggered? How does the user interact with it?
3. **Gap analysis:** Where is the biggest gap in your design — a moment where the user cannot see, understand, or override the system's intelligent behavior? What would you add to close that gap?

Submit on Canvas.

This reflection helps you strengthen your A2 before submission. Use it.

Before Next Week

— **Finish and submit A2** by Monday, April 20

— **Optional reading:** Explore any design system documentation — Material Design, Apple HIG, Shopify Polaris — look at how they handle states, feedback, and error patterns

— **Start thinking about A3:** What ethical tension exists in your A2 system? Where could your system cause harm? Who is most vulnerable?

— Next week's lecture covers **ethics, fairness, and responsible AI design** — the foundation for A3

A3 builds on A2. The ethical tensions you notice now will become the core of your final assignment.

References

Design Systems:

— Brad Frost, Atomic Design (2016) — free online

— Alla Kholmatova, Design Systems (Smashing Magazine, 2017)

— Salesforce UX, "Living Design System" (2014) — design tokens

— Dan Saffer, Microinteractions (O'Reilly, 2013)

— Amershi et al. (2019). "Guidelines for Human-AI Interaction" — CHI '19

— Weisz et al. (2024). "Design Principles for Generative AI Applications" — CHI '24

Course Frameworks:

— Parasuraman et al. (2000). "A Model for Types and Levels of Human Interaction with Automation" — IEEE SMC

— Horvitz (1999). "Principles of Mixed-Initiative User Interfaces" — CHI '99

— Johnson et al. (2014). "Coactive Design" — JCEDM

— Dey (2001). "Understanding and Using Context" — Personal and Ubiquitous Computing

— Jiang et al. (2023). "A Situation Awareness Perspective on Human-AI Interaction" — IJHCI

Media Sources

Material Design | Apple Human Interface Guidelines | Shopify Polaris | ChatGPT | GitHub Copilot | Figma AI