Ambient Cognitive Design : Designing the World Without Interfaces

A Design Governance Approach for Implicit Interaction in AI-Mediated Systems

Kirti Kamble
https://orcid.org/0009-0003-2587-6789
Email: kirti@kirtikamble.com

Abstract

Interfaces have historically functioned as compensatory structures for systems that lacked perception, context awareness, and adaptive behavior. As computation becomes increasingly perceptive and generative, interaction is no longer reliably initiated through explicit user input. Instead, cognitive effort is delegated implicitly through inference, prediction, and context-sensitive system behavior.

This condition is not emerging; it is already deployed across voice assistants, wearable devices, spatial computing platforms, and AI-mediated services. The central challenge for design is therefore no longer interface optimization, but **governance**: determining when systems may act without explicit instruction, when interaction should surface temporarily, and how human agency remains intact under uncertainty and power asymmetry.

This paper formalizes **Ambient Cognitive Design (ACD)** as a **design governance approach** for implicit interaction in AI-mediated systems. Extending affordance thinking

(Norman, 1988) into generative, context-aware environments, it introduces practical constructs—including generative affordances, manifestation thresholds, dissolution, contestability, and non-delegable cognition—that enable designers to manage implicit cognitive delegation responsibly. The contribution is grounded in long-term design practice and positioned to remain durable across evolving AI substrates and interaction modalities.

Position & Intent

This work is written from the perspective of a senior design practitioner with nearly two decades of experience designing systems across desktop, mobile, voice-based, wearable, and spatial interaction paradigms. During this period, interfaces have progressively shrunk, dissolved, and shifted from foreground tools to ambient infrastructure.

The paper does not propose a theory of intelligence, nor does it attempt to predict the future capabilities of artificial intelligence. Instead, it treats uncertainty as a permanent design condition. Its aim is to formalize a **durable design approach** for governing how interaction emerges when systems can perceive context, infer intent, and act without explicit user initiation.
In this respect, the work is intentionally aligned with Don Norman's articulation of affordances (Norman, 1988): not as a predictive theory, but as a practice-shaping lens that helps designers reason about responsibility, interpretation, and control as technological substrates evolve.

Methodological Note

This work adopts a **practice-led, analytical methodology** grounded in long-term professional experience rather than controlled laboratory experimentation. The framework presented here emerged through reflective synthesis across nearly two decades of designing and reviewing interactive systems spanning graphical interfaces, mobile computing, voice assistants, wearable devices, spatial computing, and AI-mediated services.

The approach combines:

- critical analysis of deployed systems and failure patterns,
- comparative evaluation of existing interaction paradigms,
- and theoretical grounding in established work within human–computer interaction, cognitive science, and design research.

The goal is not empirical generalization, but **conceptual consolidation**: formalizing a design governance approach that remains robust across changing technological substrates. This methodological stance aligns with precedent-setting contributions in HCI and design theory that were similarly derived from reflective practice and conceptual synthesis rather than controlled experimentation (e.g., Norman, 1988; Weiser, 1996; Dourish, 2001).

Introduction — The Interface Has Already Begun to Disappear

The disappearance of interfaces is not a stylistic trend or a speculative future; it is a structural consequence of perceptive computation. Screens, menus, and interaction flows emerged as necessary intermediaries when systems were blind to context and dependent on explicit instruction. As sensing, machine learning, and generative systems mature, the cost of explicit interaction increasingly exceeds its cognitive value.

In contemporary systems, interaction often begins **before** the user acts. Voice assistants suggest actions without prompts, navigation systems reroute without confirmation, and wearable devices intervene based on inferred physiological or situational states. In such systems, interfaces appear transiently—if at all—and frequently only after system action has already occurred.

This shift renders much of traditional UX framing insufficient. Established principles of visibility, feedback, and discoverability (Norman, 1988; Nielsen, 1994) presuppose stable interfaces and explicit user intent. In ambient and AI-mediated systems, these assumptions no longer hold. The core design problem becomes one of **governance**: deciding when systems should act implicitly, when interaction should surface temporarily, and how humans retain authorship over goals, values, and outcomes.

Why This Work Is Necessary Now

The conditions that make Ambient Cognitive Design necessary have already materialized. Large language models, perceptive sensors, and generative systems are no longer experimental; they are embedded in everyday tools, platforms, and services. Interaction increasingly occurs without explicit user initiation, while cognitive work is delegated implicitly through inference and prediction.

At the same time, existing design paradigms remain anchored to explicit interaction, persistent interfaces, and user-initiated control. This gap between technological capability and design governance has produced a growing pattern of failures—loss of trust, autonomy erosion, and unintended behavioral influence—across contemporary ambient systems.

This work arrives at a moment when design must catch up to reality. It does not speculate about future intelligence; it responds to **current deployment conditions**. Ambient Cognitive Design provides a framework for governing systems that already exist, rather than those that might one day emerge.

Figure 1: Cognitive Load Distribution in Interface-Driven vs. Ambient Systems

Interface-centric systems impose repeated and prolonged cognitive load, leading to accumulated cognitive debt over time. In contrast, governed, ambient systems minimize cognitive friction by governing when and how interaction manifests, maintaing cognitive cognitive sense.

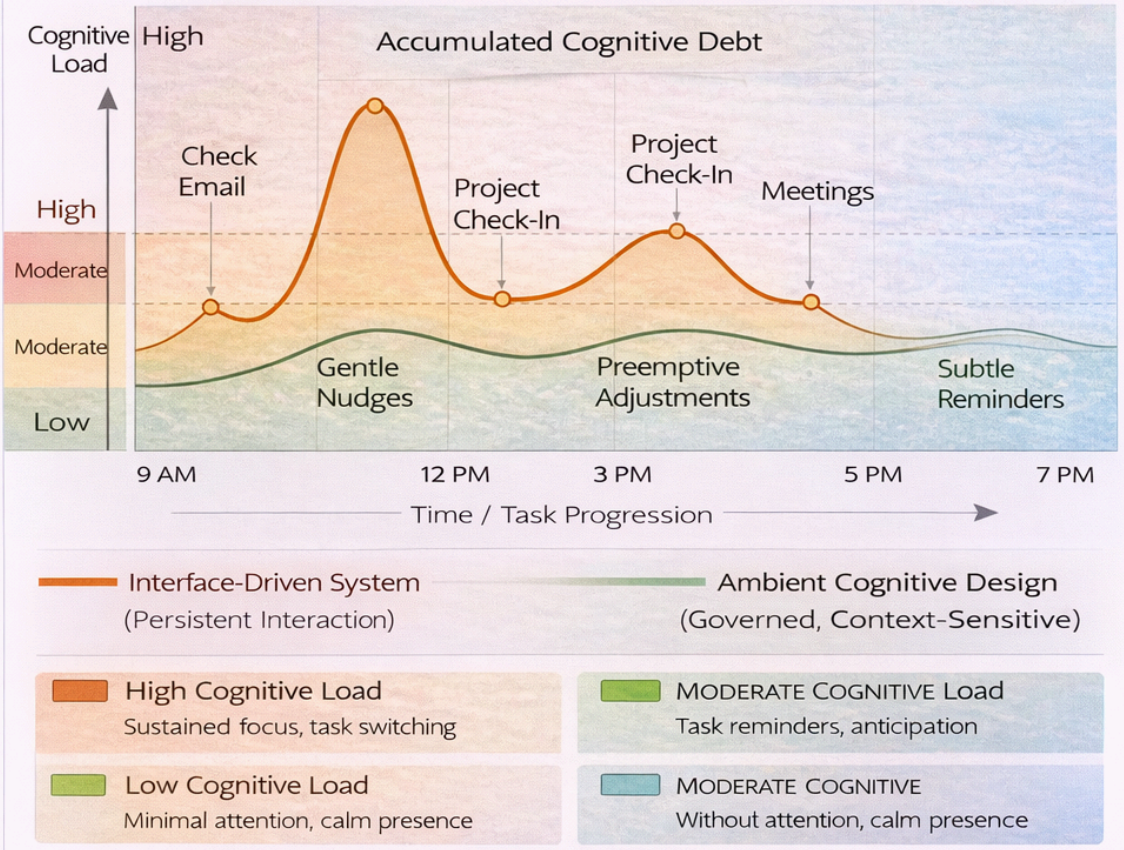

Figure 1: Cognitive Load Distribution in Interface-Driven vs. Ambient Systems.

Interface-driven systems impose repeated and prolonged cognitive load, leading to accumulated cognitive debt over time. In contrast, governed, ambient systems minimize cognitive friction by dynamically surfacing contextually relevant affordances only when situationally beneficial.

Conceptual Hierarchy: Clarifying Terms

To avoid ambiguity, this paper distinguishes three related but distinct layers:

Ambient Cognitive Design (ACD)

A **design governance approach** concerned with how implicit cognitive delegation is authorized, constrained, and contested. ACD operates at the level of system behavior, thresholds, and responsibility rather than interface form.

Ambient UX

The **design practice** that applies ACD principles to shape how interaction emerges, feels, and dissolves in lived experience. Ambient UX concerns experiential quality but is guided by governance decisions made upstream.

Ambient UI

The **temporary artifacts**—visual, spatial, auditory, or haptic—through which interaction manifests. Ambient UI elements are expressions of system behavior, not its foundation. This hierarchy reflects a deliberate inversion of traditional UX thinking: experience follows governance, not the reverse.

Part I — Why Existing Paradigms No Longer Hold

UX as a History of Cognitive Offloading

User experience design can be understood as a continuous effort to redistribute cognitive effort between humans and artifacts. Early command-line interfaces externalized memory and precision; graphical user interfaces externalized spatial reasoning; mobile interfaces externalized navigation and attention management.

What has changed is not the objective of reducing cognitive load, but the **mechanism** by which it occurs. Historically, cognitive offloading was explicit and user-initiated. In ambient systems, offloading becomes implicit and system-initiated. This transition aligns with accounts of distributed cognition (Hutchins, 1995) and the extended mind (Clark & Chalmers, 1998), but introduces new responsibilities once systems participate actively in cognitive processes.

The Visibility–Calm Tension Revisited

Weiser’s vision of calm computing emphasized technologies that recede into the background, allowing users to focus on primary activities (Weiser & Brown, 1996). However, calmness without governance risks opacity. As systems act implicitly, the absence of visible interaction can obscure causality, intent, and accountability.

The traditional tension between visibility and calm is therefore no longer a design trade-off to be optimized; it is a constraint to be governed. Designers must decide when interaction should surface to restore understanding and when silence preserves cognitive and social integrity.

Why Prior Extensions Stop Short

Situated action (Suchman, 1987), embodied interaction (Dourish, 2001), and ecological approaches to perception (Gibson, 1979) accurately describe interaction as contextual, emergent, and embodied. However, these frameworks remain largely descriptive. They do not prescribe how systems should behave once they gain the capacity to infer intent and act autonomously.

Ambient Cognitive Design addresses this gap by shifting focus from describing interaction to **governing participation** in human cognition.

Figure 2: Inversion of Interaction Design from Interface to Governance

Traditional UX design optimizes interfaces for user experience. Ambient Cognitive Design inverts the paradigm, prioritizing cognitive governance constracts.

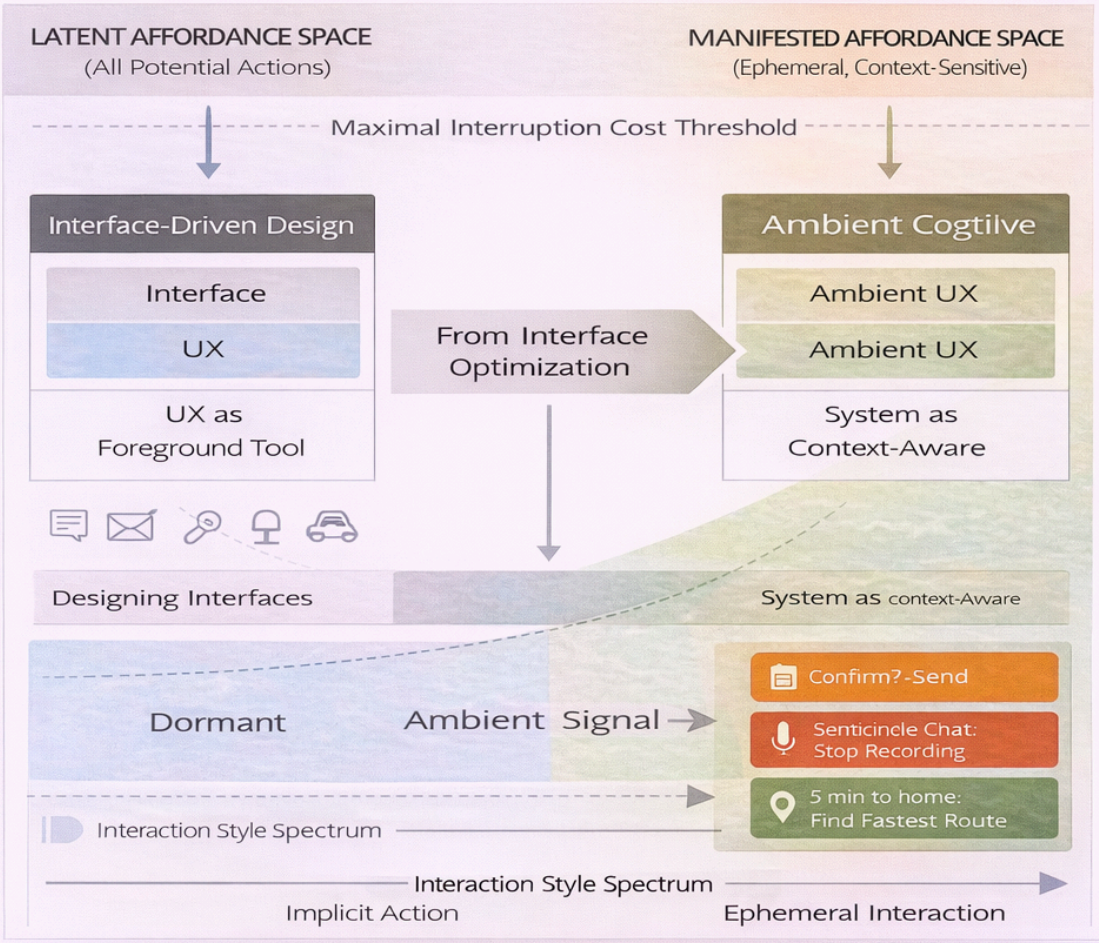

Figure 2: Inversion of Interaction Design from Interface to Governance.

Traditional interaction design centers on optimizing user interfaces for experience. Ambient Cognitive Design shifts the focus to cognitive governance, establishing rules that determine when and how interaction manifests.

PART II — GOVERNING IMPLICIT INTERACTION

Ambient Cognitive Design as a Design Governance Approach

Ambient Cognitive Design (ACD) is not a new interaction paradigm, interface style, or automation strategy. It is a **design governance approach** concerned with how and when systems participate in human cognition.

Traditional UX methodologies assume that interaction begins with explicit user intent expressed through an interface. In contrast, ambient and AI-mediated systems increasingly act on inferred intent, contextual signals, and probabilistic models. This shifts the designer's responsibility from shaping interaction flows to **governing delegation**: determining under what conditions a system may act without being asked, how its actions are made intelligible, and how control is reclaimed when inference fails.

This reframing aligns with prior observations that computation is no longer a passive tool but an active participant in situated activity (Suchman, 1987; Dourish, 2001). What ACD adds is a **prescriptive layer**: a method for constraining participation so that cognitive offloading remains supportive rather than substitutive.

Generative Affordances: Extending Affordance Thinking Beyond Static Form

Affordances, as articulated by Norman (1988), describe the actionable properties of artifacts as perceived by users. In physical and graphical interfaces, affordances are relatively stable: a handle affords pulling, a button affords pressing. Even in digital systems, affordances are designed as persistent cues.

In ambient and AI-mediated systems, affordances become **generative** rather than static. They are not pre-designed artifacts but **computed proposals** that emerge from context, inference, and probabilistic assessment. A navigation suggestion, a contextual reminder, or a spatial annotation is not always present; it manifests only when the system infers that it will reduce cognitive effort.

Generative affordances can therefore be defined as:

> *Contextually computed action possibilities that surface only when their expected cognitive benefit exceeds their attentional, social, or interpretive cost.*

This extension preserves the spirit of affordance theory while acknowledging a fundamental shift: affordances are no longer properties of objects alone, but properties of **situations** (Gibson, 1979).

Crucially, generative affordances are fragile. When surfaced prematurely or inaccurately, they cease to feel supportive and instead appear intrusive, opaque, or manipulative. This fragility necessitates explicit governance mechanisms.

Manifestation and Dissolution as Primary Design Acts

In traditional UX, the presence of interface elements is assumed. In ambient systems, **presence must be justified**.

ACD treats manifestation—the moment when interaction surfaces—and dissolution—the moment it disappears—as primary design decisions. Interaction is considered a scarce cognitive resource, not a default state.

This perspective aligns with research on attention as a limited capacity system (Kahneman, 1973) and with Weiser's principle that technology should move between center and periphery appropriately (Weiser & Brown, 1996). However, ACD makes this movement **explicit and governable**.

Manifestation Utility Principle

Interaction should manifest only when:

- the system's confidence in inferred intent exceeds a defined threshold, and
- the expected reduction in cognitive effort outweighs the cost of interruption, interpretation, or social exposure.

Likewise, interaction should dissolve once its purpose has been served. Persistent ambient UI is treated as a failure mode, not a feature.
This principle reframes interface design from layout and navigation to **temporal and situational logic**.

Implicit Cognitive Delegation: From Assistance to Substitution

Cognitive delegation refers to the transfer of cognitive work—remembering, deciding, monitoring—from humans to artifacts. Historically, this delegation was explicit: users chose to rely on calculators, calendars, or navigation tools.
In ambient systems, delegation becomes **implicit**. Systems infer when to act and often do so without a discrete moment of consent. This aligns with predictive processing accounts of cognition, where both biological and artificial systems operate by minimizing prediction error (Friston, 2010).

The design risk emerges when delegation shifts from **assistance** to **substitution**—when systems begin to act not as supports to human goals but as de facto decision-makers.
ACD therefore introduces the concept of **non-delegable cognition**: domains of decision-making that must remain under explicit human control due to ethical, social, or identity-related implications. Defining these boundaries is a design responsibility, not a policy afterthought.

Failure Patterns in Contemporary Ambient Systems

Recent attempts to commercialize ambient and AI-first devices illustrate the consequences of unguided implicit delegation.

Opaque Intent Inference (Humane AI Pin)

Early analyses of the Humane AI Pin highlighted user confusion regarding system state, intent, and action initiation. The absence of stable affordances combined with delayed or ambiguous feedback led to a breakdown in causal understanding: users were unsure why the system acted, when it was listening, or how to intervene.

From an ACD perspective, this represents a failure of **contestability** and **manifestation governance**. Interaction surfaced without sufficient intelligibility, violating Norman's principle that systems should make their state visible (Norman, 1988).

Premature Delegation (Rabbit R1)

The Rabbit R1 positioned itself as an "app-less" personal agent capable of executing tasks across services. However, early user reports indicated mismatches between inferred intent and actual goals, resulting in incorrect or partial task execution.

This reflects a breakdown in delegation thresholds. The system acted on weak signals, effectively substituting for user judgment rather than assisting it. From an ACD standpoint, this is a case of **delegation creep**, where system autonomy expands without corresponding increases in governance and reversibility.

Persistent Ambient Noise (Voice Assistants)

Longitudinal studies of voice assistants have documented declining trust and reduced usage over time, often attributed to false activations, misinterpretations, and social awkwardness (Cowan et al., 2017). These failures are not primarily technical; they are governance failures. Systems manifest too often, inappropriately, or without clear exit paths.

Contestability and Reversibility as Structural Requirements

Across these failures, a common pattern emerges: systems act without providing sufficient mechanisms for users to question, override, or recalibrate behavior. ACD treats **contestability** and **reversibility** as structural requirements rather than optional features.

Every implicit action must be:

- explainable ("Why did this happen?"),
- interruptible ("Stop"), and
- reversible within a reasonable cost window.

Without these properties, implicit interaction erodes trust even when technically accurate.

Summary

1. **Affordances in ambient systems are generative and situational**, not static artifacts.

2. **Manifestation and dissolution are the primary design acts** governing cognitive load.
3. **Implicit cognitive delegation is inevitable**, but ungoverned delegation leads to loss of agency, trust erosion, and institutional risk.

These claims shift the locus of design from interfaces to **decision rights**, setting the stage for Part III's deeper examination of agency, power, and ethical constraint.

Figure 3: Generative Affordance Computation — From Latent Possiblity to Epheneral Action

Ambient systems dynamically compute latent affordances from real-time context, manifest proposals when situationally beneficial, and dissolve them back to the latent state if not explicitly affirmed.

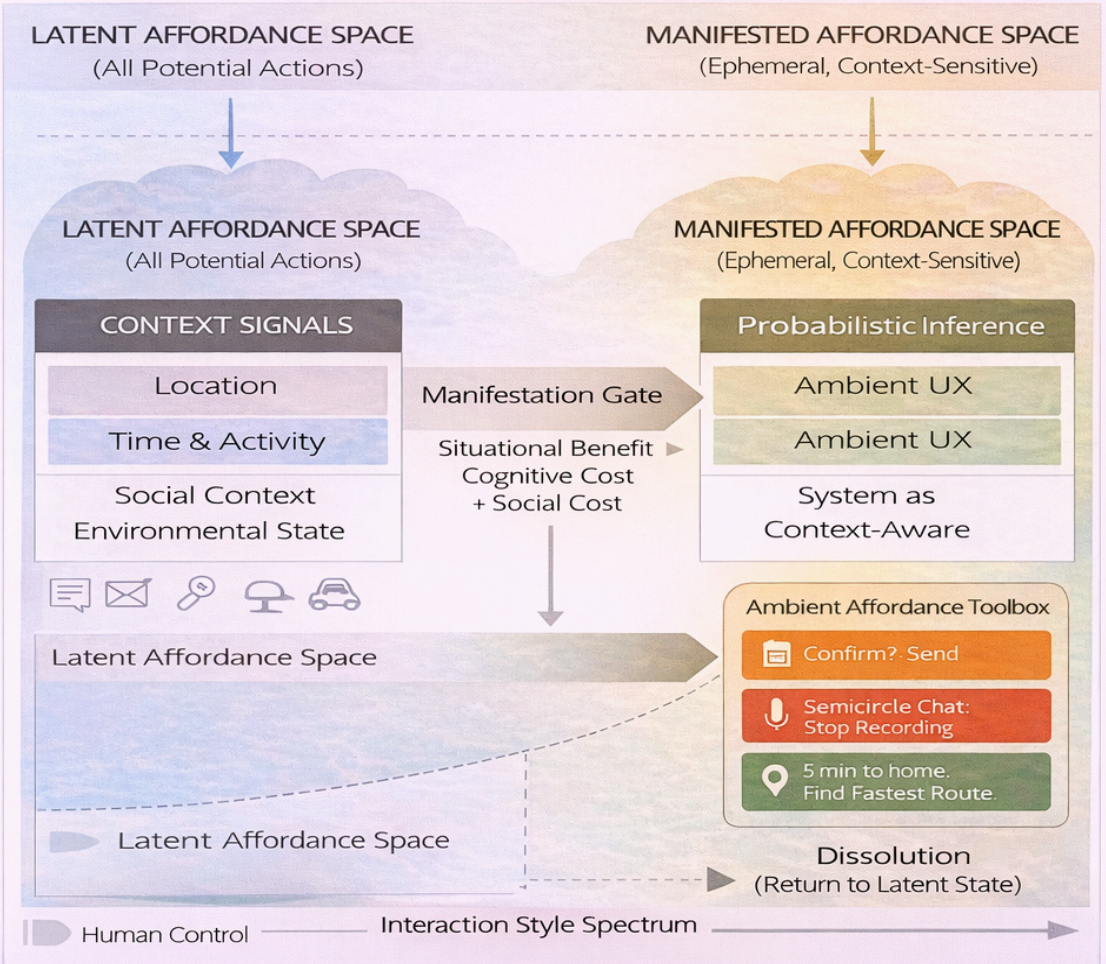

Figure 3: Generative Affordance Computation.

Ephemeral affordance computation shifts potential actions, from latent possibilities to context sensitive proposals. Each affordance is dynamically generated based on situational context, surfacing opportunistically, and then dissolves back to latent state if not explicitly affirmed.

PART III — AGENCY, POWER, AND GOVERNANCE IN AMBIENT SYSTEMS

Agency Is Not a UX Value — It Is a Structural Invariant

In traditional UX discourse, agency is often treated as a qualitative value: users should "feel in control." In ambient and AI-mediated systems, this framing is insufficient. When systems act implicitly, agency must be treated as a **structural invariant**—a condition that must hold regardless of interface presence, system accuracy, or user preference.

Agency in this context does not mean continuous manual control. Rather, it means that humans retain authorship over goals, values, and identity-defining decisions, even as systems infer intent and act on their behalf. This distinction aligns with Clark's framing of tools as extensions of human intention rather than replacements for it (Clark, 2008).

Ambient Cognitive Design therefore reframes agency as a **design obligation**, not a negotiable outcome. If agency cannot be preserved under certain conditions, those conditions must be explicitly constrained or excluded from implicit delegation.

The Power of Silence, Defaults, and Thresholds

In ambient systems, power is rarely exercised through explicit commands. It is exercised through **silence**, **defaults**, and **thresholds**.

- Silence determines when users are *not* informed.
- Defaults determine which actions occur without friction.
- Thresholds determine how much confidence is required before a system acts.

These parameters are not neutral. They encode institutional priorities, risk tolerance, and value judgments. Research on algorithmic governance has shown that seemingly technical defaults often become sites of power concentration and bias (Yeung, 2018).

From an ACD perspective, these parameters must be treated as **design artifacts** subject to review, contestation, and revision—not as implementation details.

Contestability: Designing the Right to Question

Contestability refers to a user's ability to question, challenge, or override system behavior. In explicit interfaces, contestability is often embedded through settings, confirmations, and undo actions. In ambient systems, these mechanisms are frequently absent or delayed.

Drawing on principles of procedural justice (Tyler, 2006), ACD treats contestability as essential to legitimacy. Even when system actions are correct, the absence of contestability undermines trust.

Every implicitly initiated action must therefore satisfy three criteria:

1. **Explainability** — the system can articulate why it acted.
2. **Interruptibility** — the action can be halted mid-course.
3. **Reversibility** — the effects can be undone within reasonable cost.

Without these properties, implicit interaction shifts from assistance to control.

Non-Delegable Cognition and Decision Rights

Not all cognitive work should be delegated. Some decisions are identity-defining, ethically sensitive, or socially consequential. ACD introduces the concept of **non-delegable cognition**: domains where system inference must not replace explicit human judgment. Examples include:

- moral and ethical choices,
- expressions of intent in social relationships,
- decisions with irreversible consequences,
- value-laden prioritization of goals.

This aligns with critiques of over-automation in safety-critical and social systems, where excessive delegation has been shown to reduce situational awareness and increase risk (Endsley, 1995).

Design teams must explicitly identify non-delegable domains and enforce hard boundaries around them. Treating all cognition as optimizable is a category error.

Delegation Creep and Institutional Drift

One of the most significant risks in ambient systems is **delegation creep**: the gradual expansion of system autonomy beyond its original scope. Delegation creep rarely occurs through deliberate design decisions. Instead, it emerges through incremental threshold adjustments, feature accretion, and pressure to optimize convenience.

Organizational research has shown that systems tend to drift toward automation, especially under incentives for efficiency and scale (Bainbridge, 1983). Without explicit governance, ambient systems are particularly susceptible to this drift because delegation is implicit and often invisible.

ACD responds to this risk by requiring:

- explicit documentation of delegation boundaries,
- periodic audits of system behavior,
- and design ownership of delegation decisions.

Delegation without ownership is a design failure.

Institutional Contexts and Non-Benevolent Assumptions

Much UX discourse implicitly assumes benevolent actors and aligned incentives. Ambient Cognitive Design rejects this assumption. Systems may be deployed by institutions with conflicting goals, commercial pressures, or political constraints.

Historical analyses of technological systems demonstrate that design decisions often outlive their original context and are repurposed under different power structures (Winner, 1980). Ambient systems amplify this risk because their behavior is diffuse and difficult to scrutinize.

ACD therefore insists on **designing for misuse, overreach, and capture**. Governance mechanisms must be robust enough to withstand incentive misalignment and organizational drift.

Measuring Agency and Governance Over Time

Traditional UX metrics—task success, satisfaction, efficiency—are insufficient for evaluating ambient systems. ACD proposes longitudinal indicators focused on governance health, including:

- frequency of contestation,
- reversibility usage rates,
- stability of delegation boundaries,
- trust decay or recovery over time.

These metrics shift evaluation from short-term usability to **long-term alignment**.

Summary

1. Agency in ambient systems is a **structural requirement**, not a subjective feeling.
2. Power is exercised through silence, defaults, and thresholds.
3. Contestability and reversibility are essential for legitimacy.
4. Some cognition must remain non-delegable.
5. Institutional drift is inevitable without explicit governance.

Together, these claims reposition design as a discipline responsible not just for experience quality, but for **cognitive sovereignty** in AI-mediated environments.

Figure 4: Manifestation Scarcity & Ephemerality—Governing When Interaction Surfaces

Context-sensitive systems surface ephemeral interfaces only when context relevance is high, dissolving them when no longer beneficial, minimizing interaction time and reducing cognitive load.

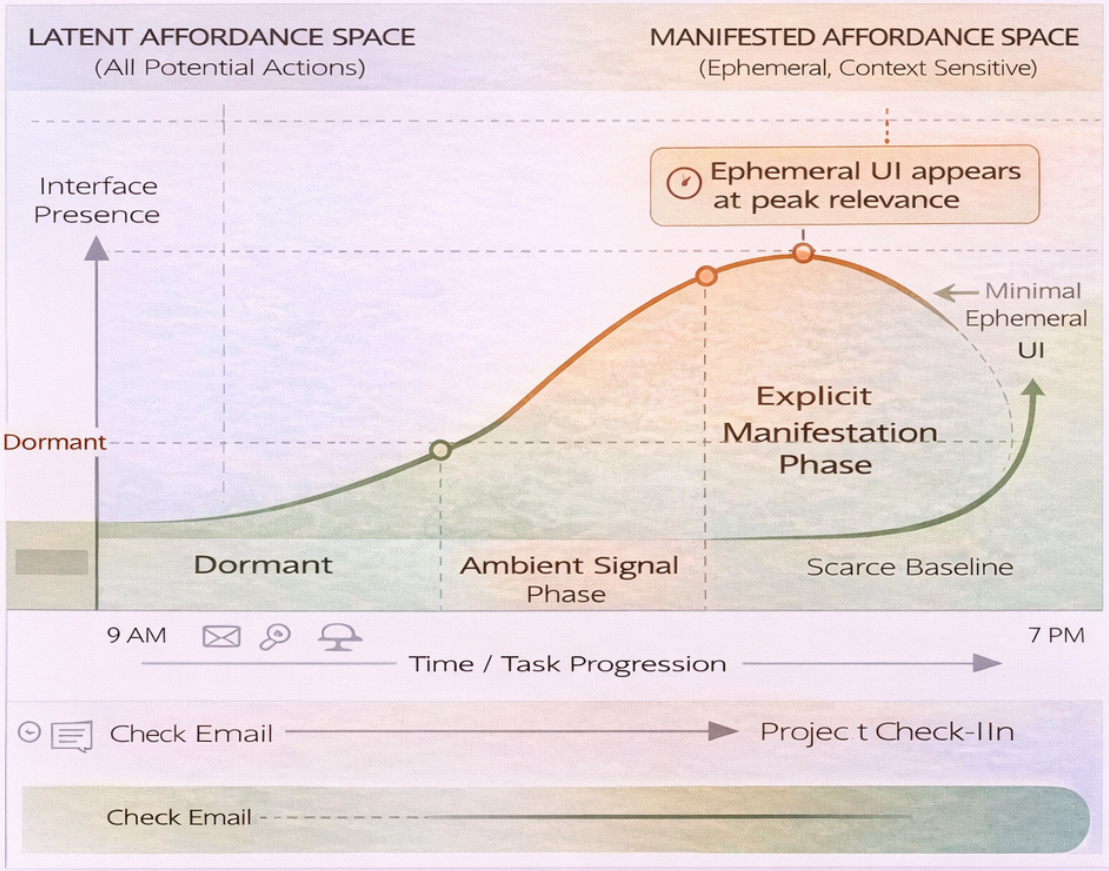

Figure 4: Manifestation Scarcity & Ephemerality

Interaction scarcity is governed temporally. Systems balance minimal presence with maximal relevance by manifesting ephemeral interfaces at peak context utility, dissolving them quickly as relevance fades.

PART IV — PRACTICE, ADOPTION, AND ORGANIZATIONAL REALITY

From UI-Centric Design to Cognitive Governance

Most contemporary organizations are structurally optimized for UI-centric design. Teams are organized around screens, features, and interaction flows; success is measured through short-term engagement and task completion metrics. Ambient Cognitive Design (ACD) does not reject these practices outright, but it reframes them as **insufficient** for systems that act implicitly.

Transitioning to ACD requires recognizing that governance decisions—thresholds, defaults, delegation boundaries—are now first-order design artifacts. These decisions cannot be retrofitted after implementation; they must be specified alongside functional requirements. This mirrors shifts observed in safety-critical domains, where human–automation interaction moved from interface optimization to responsibility allocation (Endsley & Kiris, 1995). ACD extends this logic into everyday, consumer-facing systems.

Incremental Adoption: Hybrid Interaction as the Transitional State

Despite narratives of “interface-less” futures, most real-world systems will remain **hybrid** for the foreseeable future. Explicit interfaces and implicit behaviors will coexist, often within the same user journey.

ACD treats hybrid interaction not as a failure, but as a necessary transitional state. Explicit UI serves three critical roles during this transition:

1. **Calibration** — helping users understand system behavior.
2. **Recovery** — providing a fallback when inference fails.
3. **Boundary Reinforcement** — making non-delegable cognition visible.

Research on automation adoption consistently shows that premature removal of explicit controls increases error and distrust (Parasuraman & Riley, 1997). ACD therefore recommends deliberate overlap between explicit and ambient interaction during early deployment phases.

Organizational Roles and Ownership

One of the most common failure modes in ambient systems is unclear ownership of governance decisions. Thresholds drift, defaults change, and delegation expands without a clear accountable role.

ACD requires organizations to assign explicit ownership over:

- delegation boundaries,

- manifestation thresholds,
- contestability mechanisms,
- and reversal costs.

This responsibility cannot be absorbed entirely by engineering or policy teams. It is a **design leadership function**, analogous to the role design plays in safety and ethics reviews. Studies of socio-technical systems emphasize that accountability gaps, not technical errors, are the primary drivers of systemic failure (Perrow, 1984). ACD addresses this by making governance visible and assignable.

Metrics Beyond Usability: Evaluating Governance Health

Traditional UX metrics—efficiency, satisfaction, task success—capture only short-term outcomes. Ambient systems require **longitudinal evaluation** focused on governance health. ACD proposes complementary indicators, including:

- **Contestability frequency** — how often users question or override system actions.
- **Reversibility usage** — how often actions are undone and at what cost.
- **Delegation drift** — changes in autonomy over time.
- **Trust stability** — recovery after system failure.

These metrics align with research showing that trust in automation is dynamic and fragile (Lee & See, 2004). Measuring trust recovery is often more informative than measuring initial acceptance.

Product Strategy Implications for XR and Wearable Platforms

Spatial computing platforms and wearable devices—such as head-mounted displays and ambient audio systems—are particularly sensitive to governance failures. Their proximity to the body and social context amplifies the cost of misaligned interaction.
In XR environments, persistent UI elements can obscure situational awareness, while premature system action can create safety risks. Studies in augmented reality highlight the importance of context-aware interruption management and minimal visual clutter (Billinghurst et al., 2015).

ACD reframes platform success criteria away from feature richness toward **restraint, timing, and reversibility**. This perspective aligns with emerging industry observations that adoption hinges less on capability and more on trust and social acceptability.

Institutional Deployment and the Risk of Overreach

When ambient systems are deployed at institutional scale—by governments, employers, or platforms—the risks of overreach increase. Implicit interaction can become a mechanism of surveillance, behavioral nudging, or coercion without explicit user awareness.

Critical scholarship on algorithmic governance has shown that opacity and automation can entrench power asymmetries (Eubanks, 2018). ACD directly confronts this risk by requiring explicit design constraints on delegation, explanation, and reversibility.
Importantly, ACD does not assume benevolent deployment contexts. Its governance mechanisms are designed to remain robust under misaligned incentives and regulatory pressure.

Design Education and Skill Shifts

The shift from interface design to cognitive governance has implications for design education and professional development. Traditional training emphasizes visual hierarchy, interaction flows, and usability testing. While these skills remain relevant, they are no longer sufficient. ACD demands additional competencies, including:

- reasoning about probabilistic systems,
- understanding cognitive load and attention,
- designing for uncertainty and failure,
- and negotiating ethical boundaries with technical teams.

This echoes calls within HCI education to broaden design literacy beyond artifacts to systems and consequences (Friedman & Hendry, 2019).

Summary

Part IV demonstrates that Ambient Cognitive Design is not aspirational—it is deployable under real organizational constraints. It shows that:

1. Adoption is incremental and hybrid.
2. Governance requires explicit ownership.
3. Evaluation must be longitudinal.
4. XR and ambient platforms amplify governance risks.
5. Institutional deployment demands non-benevolent assumptions.
6. Design education must evolve accordingly.

Together, these insights position ACD as a **practice-ready design governance approach**, not a speculative framework.

Part V — Limitations, Blind Spots, and Application to Contemporary Ambient Systems

---

Scope and Deliberate Limitations of Ambient Cognitive Design

Ambient Cognitive Design (ACD) is intentionally constrained. It does not attempt to prescribe system architectures, AI models, or specific interaction techniques. Its purpose is not optimization, but **governance**.
Several limitations must therefore be stated explicitly.

ACD Does Not Predict AI Capability

ACD is substrate-agnostic by design. It does not assume:

- reliable intent inference,
- generalized intelligence,
- or stable world models.

As such, it cannot resolve failures arising purely from technical insufficiency. Instead, it provides a framework for **containing harm when inference is uncertain or wrong**. This limitation is deliberate, reflecting longstanding cautions against overestimating automation competence (Bainbridge, 1983).

ACD Does Not Replace Explicit Interfaces

ACD does not advocate for interface elimination. In many domains—safety-critical tasks, learning contexts, value-laden decisions—explicit interaction remains essential.
The framework governs *when implicit interaction is appropriate*, not whether explicit interaction is obsolete. Systems that prematurely remove explicit controls risk loss of calibration, recovery, and trust (Parasuraman & Riley, 1997).

ACD Assumes Institutional Pressure and Drift

ACD does not assume benevolent actors or aligned incentives. However, it cannot fully prevent misuse in contexts where governance mechanisms are deliberately bypassed or suppressed.
In such cases, ACD serves as a **diagnostic and accountability tool**, not a technical safeguard. This limitation mirrors broader critiques of ethical-by-design approaches when detached from enforcement (Winner, 1980; Eubanks, 2018).

Blind Spots and Open Questions

Despite its breadth, ACD has identifiable blind spots that warrant further research.

Collective and Multi-User Cognition

ACD primarily addresses individual cognition. In shared or collective environments—collaboration, public spaces, multi-user XR—delegation and contestability become distributed and socially negotiated. How governance scales across groups remains an open question.

Cultural Variability in Agency Expectations

Expectations of agency, silence, and intervention vary across cultures. ACD currently assumes a broadly Western, autonomy-centric model of agency. Cross-cultural validation and adaptation are necessary to avoid imposing normative assumptions.

Emotional and Affective Delegation

ACD focuses on cognitive delegation. Emotional regulation and affective computing introduce additional complexity, where delegation may alter identity and well-being in less visible ways. This remains an underexplored domain.

Applying ACD to Contemporary Ambient Platforms

ACD is most valuable when applied not as a retrofit, but as a **lens for diagnosing why current systems struggle**.

Google XR / Android XR Ecosystem

Google's XR trajectory emphasizes environmental awareness, multimodal input, and integration with everyday tasks. However, XR platforms risk over-instrumentation: persistent overlays, excessive prompts, and fragmented agent behavior.

ACD contribution:

- Provides criteria for **when spatial UI should dissolve**
- Helps define **manifestation thresholds** in safety- and socially-sensitive contexts
- Clarifies non-delegable cognition in navigation, surveillance, and attention capture

ACD bridges the gap between *capability* and *restraint* in spatial systems.

Apple Vision Pro

Vision Pro demonstrates exceptional fidelity in explicit interaction but relies heavily on visible, windowed UI metaphors. While this supports clarity and learnability, it risks anchoring spatial computing to screen-centric paradigms.

ACD contribution:

- Offers a pathway from **explicit-first to implicit-optional interaction**
- Helps determine which interactions can become generative affordances
- Preserves Apple's emphasis on user agency through contestability and reversibility

ACD reframes Vision Pro not as "the end state," but as a **calibration phase**.

Voice Assistants (Alexa, Google Assistant, Siri)

Voice assistants illustrate the dangers of unguided implicit interaction: false activations, social awkwardness, and trust erosion over time (Cowan et al., 2017).

ACD contribution:

- Diagnoses these failures as **manifestation governance failures**
- Provides principles for reducing ambient noise
- Reframes success from responsiveness to **restraint and appropriateness**

ACD explains why technical accuracy alone has not driven sustained adoption.

Rabbit R1 and Agentic Devices

Agent-first devices like Rabbit R1 aim to eliminate apps through delegation. Early failures highlight mismatches between inferred intent and user goals, resulting in incorrect task execution and loss of confidence.

ACD contribution:

- Identifies **delegation creep** and weak thresholds as root causes
- Provides a framework for defining **explicit delegation boundaries**
- Emphasizes reversibility and calibration over execution breadth

ACD bridges the gap between agent ambition and lived usability.

Humane AI Pin

Humane AI Pin sought to remove interfaces entirely, relying on ambient interaction and projection-based feedback. User confusion around state, intent, and control reflects failures in intelligibility and contestability.

ACD contribution:

- Explains why “invisible by default” is insufficient
- Clarifies the need for **conditional visibility**
- Reinforces that calm without governance produces opacity, not trust

---

Bridging the Gap: From Capability to Governance

Across these systems, a consistent pattern emerges: failures arise not from lack of capability, but from lack of **governance framing**. Systems know *what* they can do but lack clarity about *when they should*.

ACD bridges this gap by:

- Treating interaction as a **scarce cognitive resource**
- Making delegation explicit and bounded
- Elevating contestability and reversibility to first-class requirements
- Providing designers with language to resist overreach

In doing so, ACD does not slow innovation. It **makes innovation survivable**.

Summary

This section has shown that:

1. ACD is intentionally bounded and does not overclaim.
2. Its blind spots are identifiable and researchable.
3. Its value lies in diagnosing real-world failures.
4. It provides a unifying governance lens across XR, voice, and agentic systems.
5. It bridges the persistent gap between technological capability and human trust.

Part VI — Ambient Cognitive Design, Large Language Models, and Human Agency

---

Large Language Models as Cognitive Infrastructure, Not Interfaces

Contemporary large language models (LLMs) such as ChatGPT, Gemini, and Claude represent a qualitative shift in human–computer interaction. Unlike prior automation systems, LLMs do not merely execute predefined functions; they participate in sense-making, abstraction, and linguistic reasoning across domains.

However, in their current deployment, LLMs are overwhelmingly framed as **interfaces**: chat windows, prompts, conversations. This framing obscures their more consequential role as **cognitive infrastructure**—systems that increasingly mediate how users think, remember, decide, and interpret information.
From an Ambient Cognitive Design (ACD) perspective, the central question is not whether LLMs are powerful, but **how their power is situated relative to human cognition**.

LLM Capabilities Through the Lens of Cognitive Offloading

LLMs already perform several forms of cognitive offloading that were previously effortful or unavailable:

- **Externalized working memory** (summarization, recall, synthesis)
- **Cognitive compression** (reducing complex domains into manageable abstractions)
- **Perspective generation** (alternative framings, explanations)
- **Procedural scaffolding** (step-by-step reasoning support)

These functions align closely with extended cognition frameworks, in which tools become part of the cognitive loop rather than external aids (Clark & Chalmers, 1998).
However, without governance, these same capabilities can quietly shift from **supporting cognition** to **substituting judgment**, especially when outputs are treated as authoritative rather than provisional.

The Risk of Agency Erosion in LLM-Mediated Interaction

While LLMs are often described as “assistants,” their mode of operation introduces subtle risks to human agency:

- **Epistemic drift** — users defer judgment to model-generated answers
- **Goal misalignment** — models optimize for coherence rather than user intent
- **Over-trust through fluency** — linguistic confidence masks uncertainty
- **Cognitive dependency** — reduced engagement in effortful reasoning

Empirical studies have shown that automation bias increases when systems present outputs confidently, even when incorrect (Parasuraman & Riley, 1997; Dzindolet et al., 2003). LLMs amplify this effect through natural language fluency.
ACD treats these risks not as user errors, but as **design failures**.

Ambient Cognitive Design as Guarded Cognitive Offloading

Ambient Cognitive Design explicitly positions LLMs as **background cognitive partners**, not foreground decision-makers.
In this framing:

- ACD is **always on**, but rarely visible
- It operates like a **second brain**, not a second will

- It augments perception, memory, and synthesis
- It never replaces goal-setting, value judgment, or identity-defining choice

This aligns with Norman's insistence that good design amplifies human capability without obscuring responsibility (Norman, 1988).

Agency as the Non-Negotiable Boundary Condition

ACD defines **human agency** as the invariant that constrains all implicit cognitive offloading. Agency, in this context, includes:

- **Intent authorship** — humans define goals
- **Value sovereignty** — humans decide what matters
- **Judgment ownership** — humans remain accountable for outcomes
- **Override authority** — humans can always intervene

LLMs may propose, suggest, or scaffold — but they must never finalize decisions in domains where agency is essential.
This distinction mirrors Clark's differentiation between **cognitive extension** and **cognitive replacement** (Clark, 2008).

ACD vs. Agentic AI: A Critical Distinction

ACD explicitly diverges from agentic AI narratives that frame systems as autonomous actors acting on behalf of users.

Agentic Framing	ACD Framing
System acts for the user	System thinks with the user
Automation of decisions	Augmentation of cognition
Reduced user involvement	Sustained user authorship
Convenience-first	Agency-first

This distinction is critical. Systems that act *for* users risk eroding autonomy; systems that think *with* users reinforce it.

Designing LLM Interaction Under ACD

Applied to LLM-based systems (ChatGPT, Gemini, etc.), ACD implies the following design constraints:

Provisionality Over Authority

LLM outputs must be framed as **suggestions**, not conclusions. Language, tone, and interaction structure should reinforce provisionality.

Friction Where Judgment Matters

In high-stakes or value-laden contexts, ACD intentionally introduces friction to prompt human reflection rather than optimize speed.

Visibility of Uncertainty
Models must surface uncertainty, ambiguity, or alternative interpretations, preserving epistemic agency.

Continuous Calibration
Users must be able to shape how much cognitive work is delegated over time, preventing dependency and overreach.

ACD as an Extension of Human Sensing and Sense-Making

ACD can be understood as an **extension of sensory and cognitive organs**, analogous to how eyeglasses extend vision or writing extends memory.
However, unlike physical extensions, LLMs operate at the level of meaning. This amplifies both their value and their risk.
ACD ensures that this extension remains **prosthetic, not prostheticized** — enhancing capability without redefining identity or authority.

Summary: Empowerment Without Abdication

This section establishes that:

1. LLMs already function as cognitive infrastructure
2. Ungoverned cognitive offloading risks agency erosion
3. ACD provides a principled way to integrate LLMs as background cognitive partners
4. Human agency remains the invariant that constrains all delegation
5. ACD frames intelligence augmentation as empowerment, not abdication

Figure 5: Delegation, Agency & Contestability—Where Automation Must Yield to Human

Ambient intelligent systems balance automation with human control by adhering to boundaries of delegable cognition, ensuring contestability and reversibility to uphold user autonomy with systemtic cognitions.

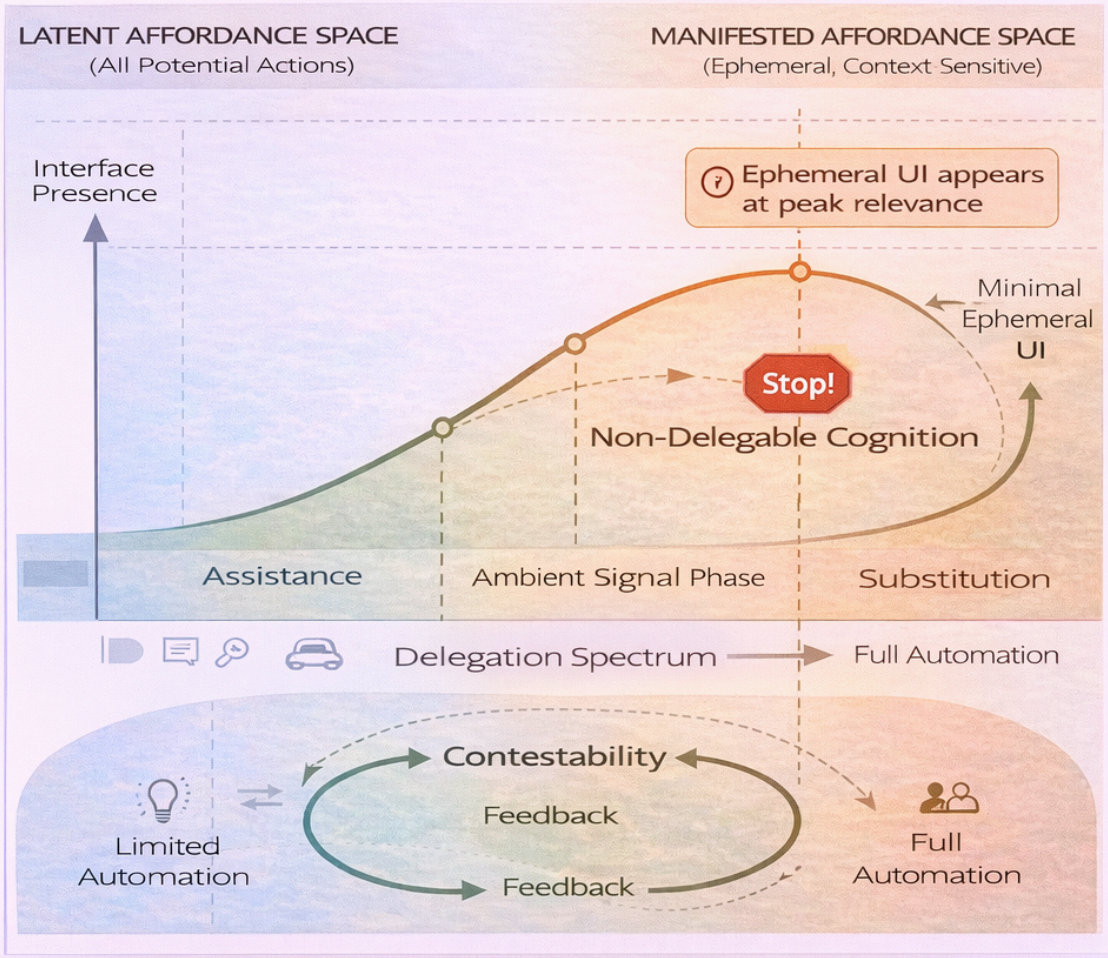

Figure 5: Delegation, Agency & Contestability

Ambient systems balance limited automation with user control. Attention-sensitive boundaries ensure that core human cognition is never delegated beyond 'Delegation'. Contestable and reversible technology preserves human agency and trust.

Conclusion — Designing Cognition, Not Interfaces

Interfaces were never the essence of interaction design. They were historical accommodations for machines that could not perceive context, infer intent, or act autonomously. As computation becomes ambient and generative, these accommodations are no longer sufficient. Interaction increasingly occurs without explicit user initiation, and cognitive effort is delegated implicitly through inference rather than instruction.

This paper has argued that this shift is not speculative. It is already underway.

The central challenge facing design is therefore no longer how to optimize interfaces, but how to **govern participation in human cognition**. Ambient Cognitive Design reframes interaction design as a discipline responsible for deciding when systems may act, when interaction should surface, and when silence preserves agency. It shifts design practice from shaping surfaces to shaping **decision rights**.

By extending affordance thinking into generative, context-aware environments, the paper introduces constructs—generative affordances, manifestation and dissolution, delegation boundaries, contestability, and non-delegable cognition—that allow designers to reason about implicit interaction under uncertainty. These constructs do not predict the future of artificial intelligence; they remain deliberately agnostic to substrate capability. Their durability lies in their grounding in human cognition, attention, and institutional reality.

Crucially, Ambient Cognitive Design rejects two false binaries that have constrained contemporary discourse. The first is the opposition between automation and control. Delegation need not imply loss of agency if boundaries are explicitly designed and contestable. The second is the opposition between calm and visibility. Interaction can remain sparse without becoming opaque, provided manifestation is governed rather than assumed. This work also makes explicit that ambient systems are not deployed in neutral contexts. Silence, defaults, and thresholds are expressions of power. Design that ignores institutional incentives, non-benevolent deployment, and long-term drift abdicates responsibility. Ambient Cognitive Design treats governance not as a policy layer added after the fact, but as a core design obligation.

The implications for practice are consequential. Designers must move beyond artifact-level decisions and accept responsibility for how cognition is redistributed between humans and systems over time. Organizations must assign ownership for delegation boundaries and evaluate success longitudinally, not merely through short-term usability metrics. Design education must expand to include reasoning about uncertainty, failure, and power.
The future of user experience is therefore not the disappearance of design, but its relocation. As interfaces dissolve, design becomes less visible but more consequential. What remains to be designed is not the screen, the gesture, or the prompt—but the conditions under which systems are allowed to think and act on our behalf.

As large language models and perceptive systems increasingly operate as continuous cognitive infrastructure, the role of design expands beyond interaction orchestration into the stewardship of human agency itself. Ambient Cognitive Design frames these systems not as autonomous agents, but as background cognitive extensions—always present, rarely foregrounded, and never authoritative over human intent. In this framing, cognitive offloading is not an abdication of responsibility, but a guarded augmentation: one that amplifies perception, memory, and sense-making while preserving authorship, judgment, and accountability.

The success of ambient AI systems will therefore not be determined by how much cognition they absorb, but by how precisely they preserve the boundary between assistance and autonomy.
Ambient Cognitive Design names this responsibility and provides a way to practice it.

Closing Position

This paper does not claim to resolve the ethical, social, or technical challenges of ambient AI. It claims something narrower and more durable: that **design has a non-negotiable role in governing cognition as computation becomes perceptive and generative**.
That role cannot be delegated.

Contributions & Implications

Primary Contributions

This paper makes **five distinct contributions** to human–computer interaction and design practice.

Reframing Interaction Design as Cognitive Governance

The paper reframes interaction design from the design of interfaces to the **governance of implicit cognitive delegation**. It argues that as systems act without explicit initiation, the central design problem shifts from usability to decision rights: when systems may act, when interaction should surface, and how agency is preserved under uncertainty.
This reframing extends prior HCI work on affordances (Norman, 1988), calm technology (Weiser & Brown, 1996), and embodied interaction (Dourish, 2001) by introducing governance as a first-class design concern.

Generative Affordances as a New Class of Affordance

The paper introduces **generative affordances**: contextually computed action possibilities that surface only when their expected cognitive benefit exceeds their attentional and social cost.
This extends affordance theory from static, artifact-bound cues to **situational, probabilistic affordances** generated at runtime. The contribution preserves the interpretive foundations of affordances while adapting them to AI-mediated systems.

Manifestation and Dissolution as Primary Design Acts

The paper formalizes **manifestation** (when interaction appears) and **dissolution** (when it disappears) as core design decisions rather than aesthetic outcomes.
This contribution provides designers with a principled way to reason about **when not to show an interface**, addressing a gap in traditional UX methods that assume persistent interaction surfaces.

Non-Delegable Cognition and Contestability as Design Invariants

The paper introduces the concept of **non-delegable cognition**—domains of decision-making that must remain under explicit human control—and frames **contestability and reversibility** as structural requirements for legitimacy.
This contribution bridges HCI with work on automation bias (Parasuraman & Riley, 1997), procedural justice (Tyler, 2006), and algorithmic governance (Yeung, 2018), translating ethical concerns into actionable design constraints. This contribution becomes particularly critical in the context of LLM-mediated systems, where fluent language generation can obscure uncertainty and accelerate automation bias.

A Practice-Ready Governance Approach for Ambient Systems

Unlike speculative or purely descriptive frameworks, Ambient Cognitive Design is articulated as a **practice-ready design approach**. It includes operational constructs, failure patterns, metrics, and organizational implications, making it applicable within real product teams and institutional settings.

Framing Large Language Models as Background Cognitive Infrastructure

The paper contributes a design framing for large language models and generative AI systems as background cognitive infrastructure rather than foreground interfaces or autonomous agents. By situating LLMs within Ambient Cognitive Design, the work provides a principled way to integrate contemporary AI capabilities as continuous cognitive support—externalizing memory, synthesis, and abstraction—without compromising human agency, intent authorship, or judgment ownership.

Implications

Implications for HCI Research

- Shifts HCI inquiry from interaction mechanics to **participation governance**
- Opens new research questions around delegation thresholds, contestability usage, and longitudinal trust
- Suggests governance-focused evaluation metrics beyond task success or satisfaction

---

Implications for Design Practice

- Requires designers to explicitly define delegation boundaries and non-delegable domains
- Elevates silence, defaults, and thresholds as intentional design artifacts
- Expands the designer’s role from experience shaping to **cognitive responsibility management**

- Designing for cognitive partnership, not delegation — Designers must treat LLMs and ambient systems as cognitive partners that support reasoning and sense-making, rather than agents that replace judgment or intent.

Implications for AI & XR Platforms

- Explains why capability-rich ambient systems fail without governance (e.g., Humane AI Pin, Rabbit R1)
- Provides a framework for designing restraint, reversibility, and calibration into XR and wearable platforms
- Positions trust and agency as adoption constraints, not secondary considerations

For LLM-powered platforms, Ambient Cognitive Design provides a counterweight to agentic narratives that prioritize task execution over cognition. It reframes success away from autonomy and toward epistemic support—helping users think better, not decide faster. This perspective has direct implications for conversational AI, multimodal assistants, and spatial computing platforms, where continuous background cognition must be carefully governed to avoid dependency, overreach, or loss of agency.

Implications for Institutions and Policy

- Demonstrates how implicit interaction can become a vector for power asymmetry and overreach
- Offers design-level mechanisms that complement regulatory and policy efforts
- Treats governance as a design obligation rather than a compliance afterthought

---

Appendices

Appendix A — Ambient Cognitive Design Checklist

Before deploying any implicit system behavior:

1. What cognition is being delegated?
2. Is this delegation assistive or substitutive?
3. What confidence threshold authorizes action?
4. How does the user contest or reverse it?
5. Is this domain explicitly delegable?
6. Who owns this decision organizationally?

If any answer is unclear, deployment is premature.

Appendix B — Common Failure Patterns in Ambient Systems

Failure Pattern	Description	Observed In
Delegation Creep	Gradual autonomy expansion	Voice assistants
Opaque Action	No intelligible causality	Humane AI Pin

Premature Inference	Acting on weak signals	Rabbit R1
Persistent Noise	Over-manifestation	Smart assistants
Irreversibility	High recovery cost	Automation workflows

Appendix C — Glossary of Key Terms

- **Ambient Cognitive Design** — A design governance approach for implicit interaction
- **Generative Affordance** — Contextually computed action possibility
- **Manifestation** — Moment interaction surfaces
- **Dissolution** — Moment interaction disappears
- **Non-Delegable Cognition** — Decisions that must remain human-controlled
- **Contestability** — Ability to question or override system action

---

Appendix D — Design Decision Matrix for Ambient Cognitive Design

The following matrix translates Ambient Cognitive Design principles into concrete decision checkpoints. It is intended as a governance aid rather than a prescriptive checklist.

Design Decision	ACD Lens	Guiding Question	Risk if Ignored
When should the system act implicitly?	Manifestation Threshold	Does acting now reduce cognitive effort more than it introduces risk or ambiguity?	Premature automation, loss of trust
When should interaction surface?	Cognitive Utility	Is visibility necessary for understanding, agency, or recovery?	Opacity, user confusion
When should interaction dissolve?	Dissolution Principle	Has the interaction served its purpose?	Persistent cognitive noise
What cognition can be delegated?	Delegable Cognition	Is this decision reversible and low-stakes?	Agency erosion
What cognition must remain human?	Non-Delegable Cognition	Does this decision define intent, values, or identity?	Loss of autonomy
How can users intervene?	Contestability	Can the user question, stop, or override this action?	Institutional overreach

How does delegation evolve over time?	Drift Monitoring	Are thresholds or autonomy levels expanding without review?	Delegation creep
Who owns this decision?	Governance Ownership	Is responsibility for this behavior clearly assigned?	Accountability gaps

---

Reference & Citation Audit

Foundational Lineage (Correctly Used)

- Norman (1988) → affordances, visibility, feedback
- Gibson (1979) → ecological perception
- Weiser & Brown (1996) → calm technology
- Suchman (1987) → situated action
- Dourish (2001) → embodied interaction

Cognitive & Automation Research

- Kahneman (1973) → attention as scarce resource
- Clark & Chalmers (1998) → extended cognition
- Hutchins (1995) → distributed cognition
- Friston (2010) → predictive processing

Governance, Power & Ethics

- Bainbridge (1983) → automation irony
- Parasuraman & Riley (1997) → misuse/disuse
- Lee & See (2004) → trust dynamics
- Yeung (2018) → algorithmic regulation
- Winner (1980) → politics of artifacts
- Eubanks (2018) → institutional harm

Empirical / Platform Context

- Cowan et al. (2017) → voice assistant trust decay
- Billinghurst et al. (2015) → AR constraints