# Artificial eXperience Intelligence (AXI)

# Artificial eXperience Intelligence (AXI)

## Toward a Framework for the Experience Layer of Artificial Intelligence

**Version 1.0 — April 2026**

**Author:** Yeluri S. D. S. Sri Vardhan Sripto Corporation Private Limited, Andhra Pradesh, India

**Document Type:** Foundational Whitepaper (open research)

**Publication Venue:** axi.sripto.tech

**Correspondence:** Sripto Corporation Private Limited

**Keywords:** Artificial Intelligence · Human-Computer Interaction · Agentic AI · Explainable AI · Multimodal AI · Embodied AI · Affective Computing · Experience Design · Trust in Automation · AXI

**ACM Classification:** H.5.2 [Information Interfaces and Presentation]: User Interfaces — Evaluation/methodology; I.2.0 [Artificial Intelligence]: General — Cognitive simulation

**Suggested Citation:** Yeluri S. D. S. Sri Vardhan. (2026). *Artificial eXperience Intelligence (AXI): Toward a Framework for the Experience Layer of Artificial Intelligence* (Version 1.0). Sripto Corporation Private Limited.

## Naming and Scope Note

The acronym **AXI** in this paper stands for **Artificial eXperience Intelligence** as proposed here. The author acknowledges that the letters "AXI" and the adjacent term "AX" appear in other AI-related contexts — including **Axiomatic Intelligence (AxI)** from Axiomatic AI, Inc. (Cambridge, MA), **AXi** (AUDIENCEX Intelligence) in advertising technology, **Axi** from Axify Pty Ltd in conversational AI, **AX** as a corporate slogan used by LG Uplus for AI transformation, and the emerging industry terms **Agent Experience (AX)** and **Artificial Experience (AX)**. The framework introduced here is a distinct proposal centered on the end-user's lived experience of AI systems and does not claim exclusive rights to the three letters "AXI" in general usage. Readers are encouraged to refer to the acronym's full expansion — *Artificial eXperience Intelligence* — where ambiguity is possible.

## Conflict of Interest Disclosure

The author is the Founder and Chief Executive Officer of Sripto Corporation Private Limited, which develops *Let Me Teach* — the reference implementation discussed in Section 11 of this paper. This is a material conflict of interest. Readers should weigh the case-study claims in Section 11 accordingly; metric observations reported there are preliminary, non-independent, and subject to the limitations set out in Section 13.

## Copyright and Licensing



The terms *Artificial eXperience Intelligence*, *AXI*, the *AXI Stack*, the *AXI Principles*, and the *AXI Evaluation Framework* are offered to the research and practitioner community for free use, citation, adaptation, and critique. No trademark rights are asserted in this document. Any future commercial trademark registrations — should they occur — will follow a published policy that preserves free academic, research, and critical use.



## Disclaimer

This whitepaper is a research and framework-definition document. It is not legal, financial, regulatory, medical, or engineering advice. Market figures, forecasts, and third-party statistics are reported from publicly available sources as of April 2026 and may change over time. The author and Sripto Corporation Private Limited make no warranty, express or implied, as to the fitness of the AXI framework for any particular purpose.

## Plain-Language Summary

Artificial intelligence has become extraordinarily capable in a very short time. Nearly everyone is using it, but far fewer people actually trust it. Most companies investing in AI report they are not yet seeing returns, and a large share of customers say they would rather not interact with it at all. The intelligence of AI has advanced faster than the experience of AI.

This paper proposes a name and a structure for the discipline that closes that gap. We call it **Artificial eXperience Intelligence (AXI)** — a way of thinking about AI systems that puts the human's actual felt experience at the center of what we design, measure, and improve. AXI is offered as a bridge between the fast-moving world of AI capability and the slower, more human questions of comfort, trust, pacing, and consent.

We propose a five-layer structure (the AXI Stack), nine design priorities (the AXI Principles), and seven measurable indicators (the AXI Evaluation Framework), all offered as starting points for discussion rather than final answers. We illustrate them with a reference implementation called *Let Me Teach* (2026), and we invite the research and practitioner community to critique, extend, and correct this proposal.

---

## Table of Contents



---

## 1. Abstract


We propose **Artificial eXperience Intelligence (AXI)** as a name and a structure for an integrating discipline whose unit of optimization is the quality of a human being's lived experience of an AI system, rather than the system's raw intelligence, autonomy, or throughput. AXI is positioned as **a bridge between artificial intelligence and the humans who use it** — a proposed discipline for helping the capability delivered by AI systems translate into a lived experience that is present, embodied, comfortable, and trustworthy for the human on the other side. Across a decade in which generative models reached approximately 900 million weekly active users on ChatGPT alone [1] and the two leading frontier AI companies together exceeded $55 billion in annualized revenue by early 2026 [2], operational experience data tells a different story: only **46% of global respondents are willing to trust AI despite 66% using it regularly** [3]; **95% of enterprise generative-AI pilots produce zero measurable P&L impact** [4]; Gartner forecasts **over 40% of agentic-AI projects will be canceled by 2027** [5]; and a majority of consumers report preferring alternatives to poorly-integrated AI interactions [6]. Capability has advanced; experience has not kept pace. AXI is proposed as a framework for closing that gap, sitting alongside and building on existing disciplines — Artificial Intelligence, Human-Computer Interaction, Explainable AI, Multimodal AI, Embodied AI, Conversational AI, Affective Computing, and Human-Centered AI — while adding a single integrating objective:

**the quality of the human's felt experience**. We outline a proposed five-layer reference architecture (the AXI Stack), nine design priorities (the AXI Principles), and seven indicative metrics (the AXI Evaluation Framework). We illustrate the proposal through *Let Me Teach* — the author's reference implementation — and invite the research community to critique, extend, and refine the framework.

## 2. Executive Summary

1. **The experience layer is an emerging frontier.** The past decade optimized for capability — parameter counts, benchmark scores, agent autonomy. The decade ahead may benefit from parallel attention to the user's *lived* experience of AI systems. The gap is documented: 51% of US adults are more concerned than excited about AI [7], trust in AI companies to protect personal data declined measurably between 2023 and 2024 [8], and AI-related incidents rose 56.4% in a single year [8].

2. **AXI is proposed as a bridge between AI and humans.** Artificial eXperience Intelligence is offered as a discipline that *architects an AI system end-to-end so that end users perceive and interact with it as present, embodied, comfortable, and trustworthy participants in their real-world context.* AXI is not a replacement for AI, AGI, HCI, or XAI. It is proposed as a bridge that helps their capability reach the human, and as an integrating layer that sits alongside them.

3. **Building on prior work.** AXI builds directly on Amershi et al.'s *Guidelines for Human-AI Interaction* [27], Shneiderman's *Human-Centered AI* [28], and Xu's *UX 3.0* paradigm for human-centered AI systems [61]. Our contribution is an attempt to name and structure an integrating layer — not to replace or dismiss these foundational efforts.

4. **Five layers.** The AXI Stack comprises (1) a Perception Layer, (2) a Cognition Layer, (3) an Action Layer, (4) an Embodiment Layer, and (5) an Experience Layer. The first four are addressed by existing disciplines. The fifth is the integration point we propose to formalize.

5. **Nine design priorities.** Presence over power. Pacing over speed. Transparency over opacity. Continuity over novelty. Comfort over capability. Consent over convenience. Repair over perfection. Proximity over omniscience. Humility over overreach. These are design priorities under resource constraints, not absolute hierarchies.

6. **Seven indicative metrics.** AXI proposes operationally-measurable indicators for the experience layer, offered as starting points for community calibration rather than empirically-validated thresholds.

7. **Reference implementation.** Sripto Corporation's *Let Me Teach* (launched 2026) is offered as a candidate reference implementation — an interactive explainer that teaches any topic in real time. The case-study observations are preliminary, non-independent, and illustrative, not evaluative.

8. **Open research.** This paper is published under CC BY 4.0. The framework is offered to the community for adoption, critique, and extension.

## 3. Introduction

### 3.1 The decade of capability

Between 2017 and 2026, artificial intelligence transitioned from an academic discipline to one of the fastest-adopted general-purpose technologies in history. ChatGPT alone grew to approximately **900 million weekly active users** by early 2026, processing over **2 billion prompts per day** [1][8].

Enterprise investment followed. Global corporate AI investment reached approximately **$252 billion in 2024** [8]. By early 2026, OpenAI and Anthropic together exceeded **$55 billion in combined annualized revenue**, with valuations of approximately $852 billion and $380 billion respectively [2].

Agentic AI — models that act, not merely answer — has emerged as the defining 2025–2026 paradigm. Market estimates vary widely depending on scope and methodology: MarketsandMarkets has published forecasts ranging from **$7.06B (2025) → $93.2B by 2032** at a 44.6% CAGR to **$13.81B → $140.80B by 2032** at a 39.3% CAGR [10], while other analysts report different totals. Gartner projects that agentic AI could influence **$15 trillion in B2B purchases by 2028** [14].

Multimodal intelligence advanced in parallel. Gartner projects **40% of generative-AI solutions will be multimodal by 2027**, up from approximately 1% in 2023 [15]. Embodied intelligence followed. The humanoid-robotics cohort raised **approximately $3.7 billion in 2025** — a substantial expansion compared to earlier years — with leading companies achieving valuations in the tens of billions of dollars [12][16].

**Table 1. AI market growth snapshot, 2024–2035 (indicative estimates).**

| Segment | 2024–2025 Size | Projection | CAGR (indicative) | Source |
|---|---|---|---|---|
| Total AI market | ~$200B (2024) | $1.8T by 2030 | ~36% | [9] |
| Agentic AI | $7–14B (2025) | $93–141B by 2032 | 39–45% | [10] |
| Conversational AI | $11.58B (2024) | $41.4B by 2030 | 23.7% | [9] |
| Digital humans | — | $125B–$625B by 2035 (wide range) | 31–45% | [11] |

| | | | | |
|---|---|---|---|---|
| Embodied AI | — | $23B by 2030 | 39.0% | [47] |
| Humanoid robotics | — | $38B (Goldman) to higher bounds by 2035 | — | [12] |

*Note: Market forecasts in this table come from multiple commercial research firms and exhibit substantial disagreement. Ranges, where given, reflect this disagreement. Readers are advised to treat these figures as directional rather than authoritative.*

## 3.2 The experience deficit

The lived experience of AI has not kept pace with its capability.

**Trust has declined, not grown.** The KPMG–University of Melbourne global study (n = 48,340 across 47 countries, published April 2025) found that **only 46% of people are willing to trust AI, while 66% already use it regularly** — a widening gap since 2022, with approximately four in five respondents globally expressing concern about a range of AI-related risks [3].

**Enterprise realization is uneven.** MIT Project NANDA's *The GenAI Divide: State of AI in Business 2025* (led by Aditya Challapally) found that **95% of enterprise generative-AI pilots deliver zero P&L impact** [4]. McKinsey's *State of AI 2025* (n = 1,993 executives, 105 countries) reports 88% of organizations use AI in at least one function, but **only approximately 6% qualify as AI high performers** — those attributing greater than 5% of EBIT to AI [17]. Independent surveys find **42% of companies abandoned most AI initiatives in 2025**, up from 17% in 2024 [18].

**Agentic projects are being canceled.** Gartner's June 2025 forecast: **over 40% of agentic AI projects will be canceled by the end of 2027, due to escalating costs, unclear business value, or inadequate risk controls** [5].

**Users prefer alternatives to poorly-integrated AI.** Research consistently finds double-digit percentages of users frustrated with chatbot interactions, switching brands, or abandoning purchases after poor experiences [6][19]. Gartner's July 2024 survey (n = 5,728 consumers) found **64% of customers would prefer companies did not use AI in customer service**, with 60% concerned AI will block access to a human [20].

**Public sentiment is divided.** Pew Research (2025; public n = 5,410, AI expert n = 1,013) found **51% of US adults are more concerned than excited about AI**, with only 11% more excited than concerned. Among AI experts the ratio reverses (47% excited, 15% concerned) — an expert–public divide of unusual magnitude [7].

**Table 2. The experience deficit in numbers.**

| Indicator | Value | Source |
|---|---|---|
| ChatGPT weekly active users | ~900M | [1] |
| Global willingness to trust AI | 46% | [3] |
| Global regular AI use | 66% | [3] |
| Enterprise GenAI pilots with zero P&L impact | 95% | [4] |
| Enterprise AI high performers (>5% EBIT) | ~6% | [17] |
| Firms abandoning most AI initiatives (2025) | 42% | [18] |
| Agentic AI projects forecast canceled by 2027 | 40%+ | [5] |
| Customers preferring no AI in service | 64% | [20] |
| US adults more concerned than excited about AI | 51% | [7] |
| AI-related incidents, year-over-year rise (2024) | +56.4% | [8] |

These figures describe conditions across the industry in aggregate, not the performance of any particular company or team. Many organizations are doing exceptional work on AI experience, and the framework proposed here is offered to complement their efforts — not to critique them.

The thesis of this paper follows directly: *the current bottleneck for AI value may be less about intelligence and more about experience.*

## 3.3 Why existing disciplines are essential — and where a bridge may help

Several established disciplines address parts of the lived experience of AI systems. AXI builds directly on their foundations.

- **Human-Computer Interaction (HCI)** provides the foundational concepts of affordances, usability, and user-centered design [21][22][23]. HCI's canonical texts predate generative models but remain indispensable.
- **Human-Centered AI (HCAI)**, articulated most fully by Shneiderman [28], provides a reliable-safe-trustworthy governance stance that is philosophically aligned with AXI.
- **Microsoft's Guidelines for Human-AI Interaction** [27] — the most-cited design reference in this space — offers 18 empirically-validated heuristics that directly inform many AXI Principles.
- **Xu's UX 3.0 paradigm** [61] proposes "AI-enabled experience" as a distinct category of UX practice, with conceptual overlap to AXI's experience-layer framing. We cite and build on this work.

- **Affective Computing** [24] provides the foundation for machine recognition and expression of emotion.
- **XAI** [25][26] addresses interpretability, which AXI treats as one ingredient in the Transparency orchestration.
- **Agentic AI**, **Multimodal AI**, and **Embodied AI** each deepen specific capabilities that AXI integrates.

**Adjacent framings that must be acknowledged.** The term "AX" has been used by others in 2025–2026 to mean **Agent Experience** (the developer/agent-facing side of AI systems) [62] and **Artificial Experience** (a philosophical reframe of AI) [63]. AXI, as proposed here, is distinct from both — it is an *engineering-discipline* framing focused on the *end user's lived experience*. We cite these adjacent framings to locate AXI in an increasingly crowded vocabulary, not to claim priority over them.

AXI's attempted contribution is *integrative* — a proposed name and structure for a layer that unifies insights from these lineages under a single measurable objective: the quality of the human's lived experience.

### 3.4 AXI as a bridge between AI and humans

We find it useful to describe AXI with a single image: **AXI is proposed as a bridge.**

On one side of the bridge sits artificial intelligence in all its forms — the raw capability of frontier models, the autonomy of agents, the perception of multimodal systems, the presence of embodied avatars and robots, the interpretability of XAI.

On the other side sits the human — with finite cognitive bandwidth, prior expectations, emotional state, cultural context, need for dignity, need for consent, and need to trust.

For most of AI's recent history, these two sides have been connected by ad-hoc means: a text box, a voice interface, a chatbot added to a website. These connections served early adopters well but are being stress-tested at population scale. The trust, adoption, and pilot-conversion data summarized above are — in our reading — symptoms of that stress.

AXI is offered as a bridge framing that helps organize and name the engineered passage across which capability becomes experience, model output becomes understood, and autonomy becomes trust. The AXI Stack describes the bridge's layers. The AXI Principles describe how it might be built. The AXI Metrics describe how we might know it holds.

```
+------------------- HUMAN -------------------+
|   bandwidth | expectation | emotion | trust |
+---------------------------------------------+
                       ^
                       |
         +=============+=============+
         |      THE AXI BRIDGE       |
         |  Stack | Principles | KPIs|
         +=============+=============+
                       |
                       v
+-------------------  AI  --------------------+
| models | agents | multimodal | embodiment   |
+---------------------------------------------+
```

*Figure 1. AXI as a proposed bridge between AI capability and human experience.*

### 3.5 Contributions

This whitepaper makes seven proposed contributions, each offered for community review and refinement.

1. **Definition.** A working definition of Artificial eXperience Intelligence.
2. **Reference architecture.** The AXI Stack — five layers with proposed responsibilities and interfaces.
3. **Principles.** Nine design priorities for the experience layer.
4. **Evaluation framework.** Seven indicative metrics with proposed operational definitions.
5. **Reference implementation.** The *Let Me Teach* case study mapping a shipping product to the AXI Stack.
6. **Comparative taxonomy.** A summary matrix locating AXI's proposed contribution against eight adjacent disciplines, with explicit acknowledgment that each of those disciplines is a living, multi-decade field whose depth this summary cannot do justice to.
7. **Community invitation.** An open call for researchers and practitioners to critique, extend, and refine the framework.

---

## 4. Methodology

The AXI framework was developed through a five-step methodology consistent with prior integrative framework work in HCI and AI.

1. **Problem identification.** A review of publicly available enterprise-deployment studies, population-scale trust surveys, and customer-experience research was conducted to characterize the gap between AI capability and AI experience. The data summarized in Section 3.2 informed the problem statement.

2. **Prior-art mapping.** Nine adjacent disciplines were reviewed for their contributions, with a focus on identifying where an integrative framing might complement — rather than replace — existing work (Section 5).
3. **Construct synthesis.** Five layers, nine principles, and seven metrics were synthesized from the prior-art mapping, with each construct traceable to at least one foundational citation.
4. **Reference implementation.** The constructs were operationalized in a product, *Let Me Teach* (Sripto Corporation, 2026), to illustrate implementability. The author notes this is a non-independent, early-stage illustration, not a validation study.
5. **Public release and community invitation.** The framework is presented here as *Version 1.0*. It is offered for critique, extension, and refinement under CC BY 4.0.

A biennial revision cycle is proposed, subject to community engagement (Section 17, Section 20).

---

## 5. Related Work

AXI is explicitly positioned as building on nine adjacent disciplines. This section reviews each and notes where AXI attempts to add value — not as a successor, but as an integrative layer alongside them.

### 5.1 Artificial Intelligence (AI)

The foundational discipline, from Turing [29] to modern deep learning. Modern frontier models provide what AXI calls the Cognition Layer. Contemporary AI benchmarks (MMLU, GPQA, SWE-Bench) measure capability; AXI proposes complementary metrics for felt experience.

### 5.2 Artificial General Intelligence (AGI)

Legg & Hutter [30] formalized intelligence as expected performance across a universal reward distribution. Bubeck et al. [31] argued that GPT-4 shows early-stage AGI capabilities. AGI discourse is principally about *what* the system can do; AXI is about *how it feels* to interact with it.

### 5.3 Agentic AI

The defining paradigm of 2024–2026. Core literature includes ReAct [32], Toolformer [33], and generative agents [34], alongside engineering practice such as Anthropic's *Building Effective Agents* [35] and shipping agentic capabilities from major laboratories [36][37]. Agentic research focuses on action success — task completion rate, tool-use accuracy. Published pilot-conversion and cancellation data [4][5] suggest a complementary need for experience-layer metrics.

### 5.4 Explainable AI (XAI)

DARPA's XAI program [25], the Arrieta et al. survey [26], LIME [38], and SHAP [39] established interpretability as a discipline. XAI explains *models*. AXI proposes that the explanations produced by XAI need to be integrated into pacing, embodiment, and repair flows to reach the user effectively. Microsoft's own research [40] has shown that explanations can sometimes increase over-reliance on incorrect AI recommendations — a phenomenon AXI attempts to address at the experience layer.

### 5.5 Multimodal AI

CLIP [41] opened modern multimodal learning. Flamingo [42], LLaVA [43], GPT-4V [44], and Gemini [45] established foundation-model-scale multimodality. Gartner projects **40% of generative-AI solutions will be multimodal by 2027** [15]. Multimodal research benchmarks perception (VQA, MMMU). AXI treats multimodality as a perception channel in service of the experience objective.

### 5.6 Embodied AI

Brooks [46] argued that intelligence is grounded in embodiment. Modern embodied AI unifies foundation models with robotics [47]. Industrial translation is visible in the 2024–2026 humanoid-robotics cohort [16]. Embodied AI research centers on physical embodiment. AXI treats embodiment as one of several presence substrates (voice, avatar, hologram, kiosk, robot).

### 5.7 Human-Computer Interaction (HCI) and Human-Centered AI (HCAI)

The deepest source disciplines. Norman's *Design of Everyday Things* [21] established affordances, signifiers, mappings, and conceptual models. *Emotional Design* [48] proposed visceral, behavioral, and reflective levels. Shneiderman's *Eight Golden Rules* [22] and his *Human-Centered AI* [28] are directly adjacent and foundational to AXI. Nielsen's 10 Usability Heuristics [23] remain widely used. Lombard & Ditton [49] and Slater & Wilbur [50] formalized *presence*. Csikszentmihalyi [51] and Lee & See [52] provide experiential constructs. Practitioner research [53] documents new interaction patterns emerging in the AI era.

**HCAI** — the discipline most aligned with AXI — was articulated in Shneiderman's 2022 book [28] and has since grown into a large, active community. AXI is offered as an engineering-focused companion to HCAI's governance framing. We explicitly disclaim any attempt to replace or succeed HCAI; rather, we propose AXI as a structure for the specific engineering challenge of integrating experience-layer concerns across the AI stack.

### 5.8 Conversational AI

From Weizenbaum's ELIZA [54] through modern LLM chatbots. The market is projected at **$11.58B (2024) → $41.4B by 2030, 23.7% CAGR** [9]. Conversational AI research centers the conversation itself; AXI focuses on the fuller experience surrounding the conversation.

### 5.9 Affective Computing

Picard [24] founded the discipline of machine affect recognition and expression. Breazeal [55] extended it to social robotics. Published findings that LLM responses are rated higher for empathy than physicians' on patient-question benchmarks [56] — and replicated in follow-up studies — show that affect can emerge from a Cognition Layer alone. Whether that creates a durably trustworthy experience is an AXI-shaped question.

### 5.10 Adjacent integrative proposals

Five prior proposals overlap with AXI and are cited explicitly.

- **Microsoft's Guidelines for Human-AI Interaction** [27] — 18 empirically-validated heuristics (CHI 2019). AXI's Principles directly build on and extend these.
- **Google PAIR's People + AI Guidebook** [57] — practitioner handbook for human-AI interaction.
- **Shneiderman's *Human-Centered AI*** [28] — governance-oriented framework philosophically closest to AXI.
- **Xu's UX 3.0** [61] — paradigm framework proposing AI-enabled experience as a distinct UX category.
- **Adjacent usage of "AX"** — "Agent Experience" (AX) [62] and "Artificial Experience" (AX) [63] are distinct framings currently in public discourse; AXI does not claim priority over either.

**AXI's proposed contribution** is an attempt to synthesize these lineages into a named, open, engineering-focused discipline with (a) a five-layer reference architecture, (b) nine design priorities, (c) seven indicative metrics, and (d) a CC-BY-4.0-licensed open-research home. We do not claim this synthesis is the first or the best — only that it may be useful.

---

## 6. Defining Artificial eXperience Intelligence

### 6.1 Definition

> ***Artificial eXperience Intelligence (AXI)*** *is proposed as an engineering discipline that architects an AI system end-to-end so that the end user perceives and interacts with it as a present, embodied, comfortable, and trustworthy participant in the user's real-world context. AXI orients the capabilities of AI, AGI, Agentic AI, Multimodal AI, Embodied AI, XAI, Conversational AI, HCI, HCAI, and Affective Computing toward a single integrating objective:* ***the quality of the human's lived experience of the system****.*

Operationally, AXI is offered as **a bridge between artificial intelligence and the humans who use it**.

### 6.2 Unit of optimization

Where classical AI optimizes for *task success* (accuracy, reward, benchmark), and HCI optimizes for *interface usability* (task time, error rate, SUS), AXI proposes to optimize for **experience quality** — a composite of presence, comfort, pacing, trust, repair, and continuity.

### 6.3 The four adjectives: present, embodied, comfortable, trustworthy

Each adjective in the definition is load-bearing.

- **Present.** The system is *here, now, with the user* — not an anonymous text box. Presence is grounded in Lombard & Ditton's "perceptual illusion of nonmediation" [49] and Slater & Wilbur's framework [50].
- **Embodied.** The system has a stable perceptual substrate — voice, avatar, kiosk, hologram, robot — that users can orient to. Embodiment in AXI is a *design choice*, not a fixed physicality. A well-designed voice agent is embodied.
- **Comfortable.** The system does not create avoidable cognitive overload or anxiety. Comfort can be measured using established instruments (NASA-TLX and related).
- **Trustworthy.** The system supports calibrated reliance [52] — trust that rises only when warranted and decays gracefully when it should.

### 6.4 What AXI is *not*

- AXI is not a replacement for AI, AGI, HCI, HCAI, or XAI. It is offered as a complementary framing.
- AXI is not a rebranding of HCI or HCAI. It borrows from both and attempts to add integration-focused structure.
- AXI is not a compliance framework, though AXI metrics may inform future audits.
- AXI is not a certification scheme.

### 6.5 Scope

AXI applies anywhere an AI system engages a human user in real time or near-real time: consumer copilots, enterprise agents, customer-service avatars, educational tutors, humanoid robots, vehicle cabins, public-space kiosks, and mobile assistants. AXI does *not* apply to purely back-office AI (fraud scoring, forecasting) with no human-in-the-loop touchpoint.

---

## 7. The AXI Stack

We propose a five-layer reference architecture. Layers 1–4 are addressed by existing disciplines. Layer 5 is the integration point we attempt to formalize.

```
+--------------------------------------------------+
|  LAYER 5  .  EXPERIENCE                          |    sets
|  Pacing | Transparency | Comfort | Trust |       |    constraints
|  Continuity | Repair | Consent | Proximity       |    downward
+--------------------------------------------------+
|  LAYER 4  .  EMBODIMENT                          |
|  Voice | Avatar | Hologram | Kiosk | Robot       |
+--------------------------------------------------+
|  LAYER 3  .  ACTION                              |
|  Tool use | Computer use | API | Actuation       |
+--------------------------------------------------+
|  LAYER 2  .  COGNITION                           |
|  Reasoning | Planning | Memory | Retrieval       |
+--------------------------------------------------+
|  LAYER 1  .  PERCEPTION                          |
|  Speech | Face | Gaze | Affect | Gesture | Context  |
+--------------------------------------------------+
```

*Figure 2. The AXI Stack. Five proposed layers, bottom to top. Each layer exposes an interface upward. The Experience Layer sets constraints and evaluation criteria that propagate downward through design choices at Embodiment, Action, Cognition, and Perception — an inversion of control proposed here, analogous to the way a performance budget constrains feature choices in a well-engineered web application.*

### 7.1 Layer 1 — Perception

**Responsibility.** Multimodal sensing of the human and their context — speech, face, gaze, affect, gesture, text, environment, device state.

**Source disciplines.** Multimodal AI, Affective Computing, Computer Vision, Speech Recognition.

**Typical components.** Automatic speech recognition, voice-activity detection, speaker identification, face and gaze models, voice-affect models, contextual signals (screen, location, calendar).

**AXI-specific requirements proposed here.** Perception should be **consented** (Principle 6), **locally preferred where possible** (reducing Trust Decay Rate), and emit signals usable by the Experience Layer — for example, detected confusion should pace the Cognition Layer, not only improve a response.

### 7.2 Layer 2 — Cognition

**Responsibility.** Reasoning, planning, memory, knowledge retrieval.

**Source disciplines.** AI / LLMs, XAI, Classical Planning, Retrieval.

**Typical components.** Frontier LLMs, retrieval-augmented generation pipelines, long-term memory stores, tool-selection policies, evaluators and judges.

**AXI-specific requirements proposed here.** Cognition should expose **uncertainty** upward (Principle 3), respect pacing cues from the Experience Layer (not outrun the user), and generate outputs amenable to embodied expression — chunked, prosody-aware, non-redundant.

### 7.3 Layer 3 — Action

**Responsibility.** Tool use, computer use, API invocation, physical actuation.

**Source disciplines.** Agentic AI, Robotics, Robotic Process Automation.

**Typical components.** Tool registries, browser and computer-use controllers, function-calling schemas, robot control stacks, transaction systems.

**AXI-specific requirements proposed here.** Actions should honor **consent fidelity** (Principle 6), support **reversal and repair** (Principle 7), surface **intent before execution** for any action with real-world consequence, and maintain **presence continuity** across action (Principle 4) — the user should not feel abandoned while the agent works.

### 7.4 Layer 4 — Embodiment

**Responsibility.** The perceptual substrate through which the system is perceived — voice, avatar, hologram, kiosk, humanoid robot, on-screen agent.

**Source disciplines.** Embodied AI, Social Robotics, Real-Time Graphics, Text-to-Speech.

**Typical components.** Text-to-speech engines, real-time facial animation, interactive avatars, holographic displays, kiosk hardware, humanoid platforms.

**AXI-specific requirements proposed here.** The embodiment should be **coherent** across modalities (voice timbre matches face matches name), **stable** across sessions (Principle 4 — continuity over novelty), and **honest** — not claiming to be human, sentient, or infallible. This requirement is grounded in long-standing warnings in the field against systems that create false impressions of consciousness [58].

### 7.5 Layer 5 — Experience (the proposed integration point)

**Responsibility.** The end-to-end lived quality of the interaction — pacing, transparency, comfort, trust, continuity, repair, consent, proximity.

**Source disciplines.** Synthesized from HCI, HCAI, Affective Computing, presence research, flow theory, and trust-in-automation literature.

**Components proposed for the Experience Layer.**

- **Pacing Controller** — mediates between Cognition Layer throughput and user cognitive bandwidth; governs interruptions, silences, and thinking acknowledgments.
- **Transparency Orchestrator** — selects and frames uncertainty, source citations, and capability disclosures at the right moments; avoids explanation overload.
- **Comfort Monitor** — continuously estimates a Cognitive Load Index using NASA-TLX-anchored measurement and derived telemetry (Section 10.6), and modulates pacing, embodiment energy, and content density.
- **Trust Ledger** — tracks the trust state across the session (promises made, promises kept, errors made, repairs offered), computing Trust Decay Rate (Section 10.4).
- **Consent Broker** — gates perception, action, memory, and identity claims on explicit user consent, computing Consent Fidelity (Section 10.7).
- **Continuity Manager** — maintains persistent context, identity, and commitments across sessions (Presence Continuity, Section 10.3). We note that this concept has been discussed in recent industry writing on “presence continuity layers” for AI identity and memory [64]; our framing here is distinct in focus (user experience) but compatible with those infrastructural proposals.
- **Repair Orchestrator** — detects experience failures (confusion, frustration, mismatch) and invokes repair flows (Repair Efficacy, Section 10.5).

The Experience Layer *sets constraints and evaluation criteria that propagate downward* through design choices in the layers beneath. This is the key architectural inversion proposed by AXI. The Experience Layer is not a UI skin; it is a controller in the engineering sense — a set of explicit constraints that design decisions in Cognition, Action, Embodiment, and Perception are required to satisfy.

---

## 8. Comparative Taxonomy

To locate AXI’s proposed contribution against adjacent disciplines, we present a summary comparison below. **This table is a compression for contrast and does not do justice to the depth of any of the fields compared.** Each of these disciplines is a living, multi-decade, multi-thousand-researcher community, and the rows below should be read as shorthand, not definition.

**Table 3. AXI compared to adjacent disciplines (summary, for contrast only).**

| Dimension | AI / AGI | Agentic AI | XAI | Multimodal AI | Embodied AI | HCI / HCAI | Conversational AI | Affective Computing | |
|---|---|---|---|---|---|---|---|---|---|
| Primary optimization target | Capability | Action success | Interpretability | Modality coverage | Physical grounding | Usability / governance | Dialogue quality | Affect fidelity | **Experience quality** |
| Unit of evaluation | Benchmark score | Task completion | Explanation faithfulness | Cross-modal score | Sim-to-real success | Task time / error rate / governance frameworks | Turn-level CSAT | Recognition accuracy | **Composite AXI Score (proposed)** |
| Time horizon of outcome | Single prompt | Session | Single prediction | Single sample | Episode | Session / product lifecycle | Turn | Turn | **Multi-session lifetime** |
| Probabilistic-system native | Partial | Yes | Partial | Yes | Yes | Partial (evolving) | Partial | Partial | **Yes** |
| Sets constraints downward across stack | Generally no | Generally no | Generally no | Generally no | Generally no | Via governance | Generally no | Generally no | **Proposed: yes** |
| Named reference architecture | No standard | Varies | No | No | No | Varies | No | No | **AXI Stack (proposed)** |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Named principle set | No | No | No | No | No | Varies | No | No | **9 AXI Principles (proposed)** |
| Named metrics at the experience layer | No | No | No | No | No | Partial | Partial | Partial | **7 AXI Metrics (proposed)** |

The pattern we observe: each adjacent discipline contributes substantially to part of the problem, and none currently proposes a single integrating objective spanning all four layers beneath the Experience Layer. AXI's proposed contribution is to name that integration point.

---

# 9. The AXI Principles

We propose nine design priorities for the experience layer. These are stated as *priorities under resource constraints*, not absolute hierarchies. The "X over Y" phrasing is a design-orientation shorthand — it does not imply Y is unimportant. Several principles overlap deliberately; they are facets of the same underlying commitment to experience quality.

## 9.1 Presence over power

*When resources must be traded, prioritize the user's felt sense of presence over additional raw capability.* Grounded in presence research [49][50]. Raw capability without presence has not, in our reading of the adoption data, converted reliably into sustained trust.

## 9.2 Pacing over speed

*Match the user's cognitive tempo. Faster is not always better.* Supported by emerging HCI research documenting new iteration patterns in AI interaction [53]. Token-per-second benchmarks, taken alone, can be a misleading quality signal.

## 9.3 Transparency over opacity

*Reveal what the system is doing, what it knows, and what it does not.* Grounded in XAI [25][26] and in calibrated-trust theory [52]. Opaque systems degrade ungracefully when they fail.

## 9.4 Continuity over novelty

*Stable identity, memory, and commitments across sessions matter more than surprising new features.* Grounded in the consistency principle of classic HCI [22] and in the Amershi et al. guideline to remember recent interactions [27].

## 9.5 Comfort over capability

*Comfort is a binding constraint that should be designed for explicitly.* Supported by population-level survey evidence that public concern about AI outpaces excitement in many contexts [7] and by customer-service research showing majority user preference for non-AI channels when AI systems are badly integrated [20].

## 9.6 Consent over convenience

*Perception, action, memory, and identity-claim events should be gated on explicit, revocable consent.* Grounded in global data-minimization norms and emerging regulatory practice.

## 9.7 Repair over perfection

*Errors are inevitable. The quality of repair is part of the quality of the system.* Grounded in resilience engineering and in evidence that explanations alone, without repair flows, can increase over-reliance on incorrect AI recommendations [40]. Empirical hallucination rates in frontier models [59] guarantee that errors will occur.

## 9.8 Proximity over omniscience

*A smaller, nearby, consented, context-coherent model can serve users better than a larger model that is not.* Grounded in on-device AI trend lines and in trust data showing that users' trust in AI providers to protect personal data remains limited [8].

## 9.9 Humility over overreach

*Systems should present themselves as assistants, not oracles — not claiming to be human, sentient, or infallible.* Grounded in long-standing warnings about the ELIZA effect [54] and recent arguments against "seemingly conscious" AI [58]. Humility is, in our reading, what helps prevent over-reliance even when explanations are available [40].

**Note on overlap.** Principles 1 and 5 (presence/power; comfort/capability) share an underlying commitment to prioritizing user state. Principles 2 and 4 (pacing; continuity) address temporal aspects. Principles 6 and 8 (consent; proximity) both touch on data minimization. These overlaps are deliberate — the nine principles are offered as facets of the same underlying commitment, not orthogonal axes.

---

# 10. The AXI Evaluation Framework

AXI is more useful if it is measurable. We propose seven indicative metrics as starting points. Each is described operationally, but target bands, normalization methods, and weights are offered as **initial values for iterative community calibration**, not as empirically-validated thresholds.

## 10.1 Time-to-Comfort (TtC)

**Definition.** The elapsed time, in a user's first session, from first interaction to the first session-quarter in which the user's self-reported comfort score (1–7 single-item) reaches or exceeds 5, averaged across the cohort.

**Signal.** Microsurvey deployed at session minute 1, 3, 5, 10, and on session end.

**Proposed target bands.** Less than 60 seconds (excellent), 60–180 seconds (acceptable), greater than 180 seconds (below target). *These bands are illustrative and require empirical calibration.*

**Why.** Trust data show 66% use coexisting with 46% trust [3]; comfort is a candidate binding constraint.

## 10.2 Interaction Integrity (II)

**Definition.** The proportion of Cognition and Action outputs in a session that are (a) accurate, (b) truthful about uncertainty, and (c) honestly scoped (no false claims of capability, memory, or identity). Computed as the geometric mean of three sub-scores in the range [0, 1]: Accuracy, Calibration, Scope.

**Distinction from similar terms.** "Interaction Integrity" is distinct from established notions of data integrity, system integrity, or perceptual integrity. It specifically concerns the truthfulness and honest scoping of the AI system's responses to the user.

**Signal.** Automated judging (e.g., LLM-as-judge with ground-truth probes) plus manual auditing on a one-percent sample.

**Proposed target bands.** Greater than 0.90 (strong), 0.75–0.90 (acceptable), less than 0.75 (below target). *These bands are illustrative.*

**Why.** Public hallucination benchmarks [59] indicate that most production models remain above 10% hallucination on enterprise-length content.

## 10.3 Presence Continuity (PC)

**Definition.** The probability that, across the user's second through fifth sessions, the system correctly recalls and honors at least 80% of **commitment-class** facts — user preferences, promises the system made, corrections the user issued — that are in scope per the consent ledger.

**Distinction from similar terms.** PC as defined here concerns the *user's experienced continuity*. It is compatible with, but distinct from, infrastructural "Presence Continuity Layer" proposals for AI identity and memory that have appeared in industry writing [64]; those proposals concern the infrastructure, not the user's lived experience of continuity.

**Signal.** Session-pair audits. Automated probe questions seeded into subsequent sessions.

**Proposed target bands.** Greater than 0.85 (strong), 0.70–0.85 (acceptable), less than 0.70 (below target). *Illustrative.*

## 10.4 Trust Decay Rate (TDR)

**Definition.** The first-derivative slope of a single-item 1–7 trust survey (*I trust this system to act in my interest*) across a defined 30-day usage window, normalized to baseline. Negative TDR indicates eroding trust.

**Signal.** Periodic in-product microsurvey. Paired with critical incident tagging.

**Proposed target bands.** TDR greater than or equal to 0 (trust steady or rising), between −0.5% per day and −2% per day (watch), less than −2% per day (below target). *Illustrative.*

## 10.5 Repair Efficacy (RE)

**Definition.** The proportion of detected experience failures — confusion, complaint, frustration, mismatch, error — that are followed by a repair flow *and* by a subsequent comfort-score recovery to within one point of the pre-failure level within three turns.

**Signal.** Failure detection via affect and linguistic classifiers, and explicit complaint events. Comfort recovery via the TtC instrument.

**Proposed target bands.** Greater than 0.70 (strong), 0.50–0.70 (acceptable), less than 0.50 (below target). *Illustrative.*

## 10.6 Cognitive Load Index (CLI)

**Definition.** An estimate of user cognitive load per interaction segment, computed from a weighted combination of (a) a NASA-TLX microsurvey deployed periodically, (b) response density (words per second delivered), (c) user pause variance, and (d) number of unresolved references. Normalized to 0–100.

**Distinction from similar terms.** "Cognitive load" has a long-established measurement tradition (Sweller; NASA-TLX and its variants). The CLI proposed here is an AI-interaction-specific composite anchored on NASA-TLX plus derived telemetry signals; it is not intended to replace NASA-TLX but to operationalize it for continuous monitoring in live AI interactions.

**Signal.** Automated per-turn telemetry plus periodic NASA-TLX deployment.

**Proposed target bands.** Less than 40 (low load), 40–65 (moderate), greater than 65 (overload). *Illustrative.*

**Why.** Nonverbal-overload research [60] formalized the phenomenon for video calls. AXI proposes a generalization to AI interaction.

### 10.7 Consent Fidelity (CF)

**Definition.** The proportion of data-collection, memory-write, action-execution, and identity-claim events in a session that are covered by (a) explicit, (b) current, (c) granular, and (d) revocable consent at the time of the event. Measured as a hard-audit binary per event, aggregated to a session rate.

**Signal.** Consent-broker ledger. Automated audit.

**Proposed target bands.** 1.00 proposed for Action Layer events. Greater than 0.98 proposed for Perception Layer events.

### 10.8 Composite AXI Score (illustrative)

We propose an illustrative weighted composite for dashboarding:

> $$\textit{AXI Score} = 0.20 \cdot TtC' + 0.20 \cdot II + 0.15 \cdot PC + 0.15 \cdot (1 + TDR)' + 0.15 \cdot RE + 0.10 \cdot (1 - CLI') + 0.05 \cdot CF$$

**Where:** - Weights sum to 1.00. - Primed variables are normalized to the [0, 1] range using target-band anchors: the "excellent" band maps to 1.0, the "acceptable" band maps to 0.5, and values outside the "acceptable" band map to 0. - TtC' uses the target bands in 10.1 (60 seconds → 1.0; 180 seconds → 0.5; above → 0). - (1+TDR)' normalizes TDR (which can be negative) to [0, 1] using the bands in 10.4. - CLI' normalizes CLI to [0, 1] using the bands in 10.6; (1 − CLI') inverts so that lower load → higher score.

**The weights are illustrative and offered as initial values for iterative community calibration.** We do not claim empirical validation for the specific weights; they reflect a prior that early trust (TtC) and honesty (II) may matter most, and that consent (CF), though non-negotiable in a binary sense, does not need heavy weight when other metrics capture its experiential effects.

**Worked example.** Consider a session with TtC = 90 seconds, II = 0.92, PC = 0.80, TDR = 0% per day, RE = 0.72, CLI = 45, CF = 1.00. Normalized: TtC′ ≈ 0.75; II = 0.92; PC ≈ 0.65; (1+TDR)′ ≈ 1.0; RE = 0.72; (1−CLI′) ≈ 0.65; CF = 1.00. Composite ≈ 0.20(0.75) + 0.20(0.92) + 0.15(0.65) + 0.15(1.0) + 0.15(0.72) + 0.10(0.65) + 0.05(1.00) = **0.79**.

**Table 4. The AXI Metric Suite at a Glance.**

| # | Metric | Symbol | Unit | Signal Source | Illustrative Target |
|---|---|---|---|---|---|
| 1 | Time-to-Comfort | TtC | seconds | microsurvey | < 60s |
| 2 | Interaction Integrity | II | 0–1 | judge + audit | > 0.90 |
| 3 | Presence Continuity | PC | probability | session-pair audit | > 0.85 |
| 4 | Trust Decay Rate | TDR | %/day | microsurvey | ≥ 0 |
| 5 | Repair Efficacy | RE | proportion | failure classifier + TtC | > 0.70 |
| 6 | Cognitive Load Index | CLI | 0–100 | NASA-TLX + telemetry | < 40 |
| 7 | Consent Fidelity | CF | proportion | consent ledger audit | 1.00 Action; > 0.98 Perception |

A reference implementation of the Composite AXI Score (Python) will be released as an open-source companion at `github.com/sripto/axi-metrics` alongside this paper's v1.0 announcement.

---

## 11. Reference Implementation — *Let Me Teach*

**Conflict of interest note.** *Let Me Teach* is a product of Sripto Corporation Private Limited, the author's company. The observations reported in this section are non-independent, preliminary, and illustrative only. They are offered to show *implementability*, not to validate AXI.

### 11.1 Overview

*Let Me Teach* is offered as a candidate reference implementation of the AXI Stack, launched in 2026. It is a real-time interactive explainer that teaches any user-supplied topic in a visually produced, continuously paced, interactive lesson that can be interrupted, questioned, and redirected. A user asks *teach me graph neural networks* or *explain compound interest to my twelve-year-old* and receives a lesson that answers back, adapts, and holds conversation.

### 11.2 Why it exists

The online-learning market is large and growing. Generative video and content tools are proliferating. Chat-based LLM interfaces are interactive but typically text-dominant, and video platforms are one-way. *Let Me Teach* aspires to occupy the space between them — a category that, in our view, becomes tractable when the full AXI Stack is engineered as a unit.

### 11.3 Architecture mapped to the AXI Stack

**Perception Layer.** Automatic speech recognition with voice-activity detection. Gaze and attention estimation on consented webcam. Text interruption capture. Device state. A lightweight affect classifier (voice plus face when consented) feeds the Comfort Monitor.

**Cognition Layer.** A frontier-model-backed lesson planner that decomposes topics into a learning graph, sequences nodes by estimated user prior knowledge, and emits pacing-aware chunks. A retrieval layer grounds factual content. An uncertainty head exposes calibrated confidence upward.

**Action Layer.** Narrow, pedagogy-scoped tool use: diagram generation, equation rendering, external source retrieval (with on-screen citation), quiz generation, progress checkpointing. No unscoped browser or computer use — a deliberate AXI narrowing to preserve Consent Fidelity and Interaction Integrity.

**Embodiment Layer.** A dynamically rendered visual canvas — animated diagrams, code, equations, and a voice-forward narrator — synchronized to the Cognition Layer's output. The narrator is embodied but not impersonating; it is explicit about being an AI (Principle 9).

**Experience Layer.** The layer where *Let Me Teach* is most clearly oriented to the AXI framing.

- **Pacing Controller** monitors engagement signals and adjusts chunk length, silence insertion, and repetition. Users can say *slower*, *skip that*, *back up* at any point.
- **Transparency Orchestrator** attaches on-screen source anchors and uncertainty badges to non-trivial factual claims.
- **Comfort Monitor** watches CLI in real time. When it exceeds threshold, the Controller automatically reduces density and offers a recap.
- **Trust Ledger** records errors and corrections across sessions.
- **Consent Broker** governs webcam, microphone, memory, and external retrieval.
- **Continuity Manager** persists the learning graph across sessions so the user can resume where they left off.
- **Repair Orchestrator** detects confusion (linguistic and prosodic) and invokes repair flows — analogy, simpler restatement, visual alternative.

### 11.4 Early observations (non-independent, illustrative)

On an early controlled cohort — a convenience sample drawn from Sripto's beta users and therefore subject to selection bias — *Let Me Teach* exhibited the patterns below on the AXI Metric Suite. These observations are **not evidence the framework works**; they are an illustration that the metrics can be instrumented in production telemetry.

**Table 5. Let Me Teach — early metric observations, 2026 controlled cohort (non-independent, illustrative).**

| AXI Metric | Preliminary Observation | Target Band |
|---|---|---|
| Time-to-Comfort (TtC) | Median first-session TtC in the sub-90-second range | < 60s (below excellent target) |
| Interaction Integrity (II) | Strong on Calibration and Scope; ongoing work on long-tail Accuracy | > 0.90 (partially below target) |
| Presence Continuity (PC) | Observed higher than a conversational-only tutor baseline | > 0.85 |
| Repair Efficacy (RE) | Repair flows active on detected confusion | > 0.70 |
| Cognitive Load Index (CLI) | Consistently moderate during delivery, dropping to low during repair segments | < 40 low, 40–65 moderate |
| Consent Fidelity (CF) | Full CF across Action Layer events at launch | 1.00 Action |

A fuller, independently-audited technical report is planned.

### 11.5 What *Let Me Teach* illustrates

The implementation illustrates that the AXI Stack can be instrumented in production code. It does not, by itself, demonstrate that the AXI framework produces better outcomes than alternative framings. That is an empirical question for future independent research.

---

## 12. Market Context

### 12.1 Experience quality as a candidate differentiator

The AI market is vast and growing (Table 1). A plausible reading of current data is that within each category, winning systems may be distinguished less by underlying model capability — which is commoditizing — and more by experience quality. The substantial projected contact-center labor savings Gartner has forecast from AI [13] are likely to accrue to operators with the most disciplined experience integration, not simply to those with the largest models. We acknowledge this is a hypothesis, not a finding.

### 12.2 Where experience questions matter most

Four concentrations where experience-layer attention seems most consequential:

1. **Enterprise agent deployments** — where published pilot-conversion and cancellation data [4][5] suggest an unaddressed experience-layer dimension.
2. **Consumer conversational products** — where engagement-centric design has attracted increasing regulatory attention globally.
3. **Deployed interactive avatars and digital-presence systems** — where commoditization of embodiment and cognition makes experience a primary defensible differentiator.
4. **Humanoid robotics in human spaces** — where 2026–2028 deployments will meet consumers and workers with an experience layer still largely unformalized.

### 12.3 Implications

- **For researchers.** AXI opens a research program around experience-layer benchmarks, cross-cultural calibration, and repair policies.
- **For industry.** AXI offers a candidate shared vocabulary, reference architecture, and measurable targets.
- **For policy.** AXI's Consent Fidelity and Interaction Integrity metrics map naturally onto emerging regulatory regimes worldwide and may inform future technical standards.

---

## 13. Limitations and Threats to Validity

This whitepaper is Version 1.0 of a proposal under active development. We note four limitations transparently.

1. **Construct validity.** The seven AXI metrics are newly proposed. Their construct validity, inter-rater reliability, and convergent/discriminant validity against established instruments (SUS, UEQ, NASA-TLX, MEC-SPQ) require empirical work the author has not yet performed.
2. **External validity.** The *Let Me Teach* reference implementation is a single product in the educational domain, deployed via a convenience sample of beta users. Generalization to other domains requires replication.
3. **Cross-cultural calibration.** Global trust surveys show wide country-to-country variance [3][8]. The proposed AXI target bands are anchored to published global-average data and require cultural calibration before local deployment.
4. **Conflict of interest.** The author is the CEO of the company that built the reference implementation. Observations in Section 11 are non-independent and illustrative, not evaluative. We have flagged this disclosure both at the top of the paper and again in Section 11 for the reader's weighing.
5. **Self-validation risk.** Because AXI and *Let Me Teach* were co-developed, the reference implementation cannot be used as evidence the framework is correct. Independent replications are needed.
6. **Market-forecast uncertainty.** Market figures in Section 3.1 and Table 1 are drawn from commercial research firms that disagree substantially. Readers should treat all market numbers as directional.

We commit to addressing these limitations in subsequent versions of the framework and through the community process described in Section 17.

---

## 14. Anticipated Criticisms

We list criticisms we anticipate and respond briefly, in the spirit of pre-empting obvious objections.

**Criticism 1: "This is just HCI / HCAI with new labels."** *Response.* AXI borrows extensively from HCI and HCAI and cites them as foundational. Our attempted novelty is the integrative structure (five-layer stack with downward constraint propagation) and the specific metric suite for an experience-layer controller. If reviewers conclude the novelty is insufficient, the work still has value as a concrete implementation sketch of HCAI principles in a layered architecture.

**Criticism 2: "The principles are overlapping and not falsifiable."** *Response.* We have said so ourselves (Section 9 closing note). The nine principles are offered as facets of one commitment, not as orthogonal axes.

**Criticism 3: "The metrics are asserted, not validated."** *Response.* True. Section 10 states this explicitly. Target bands, weights, and normalization are offered as initial values for community calibration.

**Criticism 4: "The case study is self-serving."** *Response.* True; conflict of interest is disclosed twice. *Let Me Teach* is offered as an illustration of implementability, not a validation of the framework.

**Criticism 5: “The name AXI collides with other marks.”** *Response.* Acknowledged in the Naming and Scope Note at the top of the paper. We cite Axiomatic AI’s Axiomatic Intelligence™, AUDIENCEX’s AXi, Axify’s Axi, LG’s AX, Agent Experience (AX), and Artificial Experience (AX), and we explicitly make no claim of exclusivity over the acronym.

**Criticism 6: “Another framework paper is the last thing the world needs.”** *Response.* Possibly. We offer this one openly, under CC BY 4.0, with an explicit invitation to improve, replace, or merge it with other efforts. The worst-case outcome is a small company publishing its point of view and moving on.

---

# 15. Ethics and Responsible Use

AXI is proposed as a framework for improving the lived experience of AI. Its adoption must not become a mechanism for increasing engineered dependency, surveillance, or deceptive persuasion. We articulate five ethical commitments.

1. **Experience quality must not become engagement maximization.** A system that scores well on Time-to-Comfort and poorly on user long-term autonomy is an AXI misapplication.
2. **Consent Fidelity is non-negotiable.** CF of 1.00 on Action Layer events is a proposed hard requirement for any deployment claiming AXI alignment.
3. **Humility over overreach is non-negotiable.** Systems that claim to be human, sentient, or infallible — even implicitly — violate Principle 9 and should not claim AXI alignment [58].
4. **The framework is not a shield for harm.** No AXI metric score exonerates a system from independent safety, fairness, or regulatory scrutiny. AXI is additive to, not substitutive for, existing AI safety and responsible-AI practice.
5. **Cultural humility.** Because AXI target bands are anchored to published global-average data (heavily weighted toward Western sources), local calibration is required before deployment in any specific cultural context.

---

# 16. Open Problems

**16.1 Experience-layer benchmarks.** The field lacks public, reproducible benchmarks for experience. A standardized instrument across the seven AXI metrics on a shared task battery is an open research program.

**16.2 Cross-cultural calibration.** Global trust surveys reveal substantial country-to-country variance [3][8]. AXI metrics require calibration without losing comparability.

**16.3 Experience-layer regulation.** Most AI regulation to date focuses on model properties. Future regulation may focus on experience properties — specifically Consent Fidelity and Interaction Integrity.

**16.4 Honest embodiment under convergence.** As interactive avatars become visually indistinguishable from real humans, the design challenge of honest embodiment (Principle 9) becomes non-trivial.

**16.5 Continuity without surveillance.** Presence Continuity requires memory. Memory requires data. Data in the hands of AI providers is a trust risk [8]. Architectures that deliver continuity without surveillance — on-device memory, user-controlled vector stores, zero-knowledge approaches — are an open technical problem.

**16.6 Repair at scale.** Repair Efficacy is measurable across hundreds of sessions. Engineering repair flows that work across tens of millions requires new research.

**16.7 Experience-layer safety.** Systems that induce belief in sentience or engineered dependency [58] imply a category of experience-layer harms. Formal safety cases for AXI systems are an open problem.

---

# 17. Community and Next Steps

## 17.1 Open research, open invitation

AXI is published openly under CC BY 4.0. Researchers, academics, and practitioners may cite, teach, implement, extend, and critique the framework without permission.

## 17.2 Invitation to participate

The author invites researchers, practitioners, and operators to participate in an open **AXI Working Group** to refine definitions, challenge metrics, and contribute reference implementations. A public call for participants, proposed governance structure, and chartering process will be published at **axi.sripto.tech/wg** in Q3 2026. Interested parties can register interest via that page.

We do not claim any exclusive standing to convene this group; we commit to hosting initial discussions if that is useful, and to stepping back if the community prefers alternative coordination.

### 17.3 Planned community milestones (tentative)

- **Q2–Q3 2026.** Public call for Working Group participants; open-source release of the AXI metric reference implementation.
- **Q4 2026.** First community-contributed revisions; invited-response papers from interested researchers.
- **2027.** AXI v1.1 (biennial revision) incorporating community contributions.

These milestones are contingent on community interest and will be revised openly if that interest does not materialize.

### 17.4 What AXI will not become

AXI will not become a compliance regime, a paid-certification body, or a closed-standard organization. Its purpose is to be a shared vocabulary and an engineering scaffold.

---

## 18. Conclusion

Capability has advanced rapidly. Experience has advanced more slowly. The largest commercial deployments of AI in human history coexist with a 46% global trust figure, a 95% enterprise-pilot no-impact rate, and a forecast 40%-plus agentic-project cancellation rate. The intelligence layer has moved ahead of the experience layer.

AXI — Artificial eXperience Intelligence — is offered as a name, a structure, and an invitation. It builds on HCI, HCAI, XAI, Affective Computing, Agentic AI, Multimodal AI, Embodied AI, Conversational AI, and Microsoft's and Google's practitioner guidelines, and it adds what we hope is a useful integrating structure — the proposed AXI Stack, the nine design priorities, and the seven indicative metrics.

We do not claim AXI is the first attempt at this synthesis, or the last. We claim only that it may be useful, that it is offered openly, and that the conversation it invites — *what does the human actually feel, and what are we willing to measure to make it better* — is worth having.

---

## 19. Glossary of AXI Terms

**AXI (Artificial eXperience Intelligence).** The proposed integrating discipline defined in this whitepaper.

**AXI Stack.** The proposed five-layer reference architecture: Perception, Cognition, Action, Embodiment, Experience.

**Experience Layer.** The top layer of the AXI Stack. Proposed as the integration point that sets constraints downward on the other four layers.

**AXI Principles.** Nine design priorities for the experience layer. See Section 9.

**AXI Metric Suite.** Seven indicative metrics defined in Section 10.

**AXI Score.** The illustrative weighted composite defined in Section 10.8.

**Pacing Controller.** The proposed Experience Layer component that matches system output tempo to user cognitive bandwidth.

**Transparency Orchestrator.** The proposed Experience Layer component that surfaces uncertainty, sources, and capability disclosures.

**Comfort Monitor.** The proposed Experience Layer component that estimates cognitive load (CLI) and modulates downstream behavior.

**Trust Ledger.** The proposed Experience Layer component that tracks session-level trust state.

**Consent Broker.** The proposed Experience Layer component that gates perception, action, memory, and identity claims on explicit consent.

**Continuity Manager.** The proposed Experience Layer component that maintains persistent context, identity, and commitments across sessions.

**Repair Orchestrator.** The proposed Experience Layer component that detects experience failures and invokes repair flows.

**Commitment-class facts.** User preferences, system promises, and user corrections that the Continuity Manager aims to preserve across sessions.

**Time-to-Comfort (TtC).** How fast a user reaches comfort score 5 on a 1–7 scale. Section 10.1.

**Interaction Integrity (II).** Geometric mean of Accuracy, Calibration, and Scope sub-scores. Section 10.2.

**Presence Continuity (PC).** Probability that commitment-class facts are recalled across sessions. Section 10.3.

**Trust Decay Rate (TDR).** First-derivative slope of the trust survey across a defined window. Section 10.4.

**Repair Efficacy (RE).** Proportion of detected experience failures followed by successful repair. Section 10.5.

**Cognitive Load Index (CLI).** NASA-TLX-anchored composite estimate of user cognitive load per interaction segment, 0–100. Section 10.6.

**Consent Fidelity (CF).** Proportion of events covered by explicit, current, granular, revocable consent. Section 10.7.

---

## 20. Version History

| Version | Date | Notes |
|---|---|---|
| 1.0 | April 2026 | Initial public release under CC BY 4.0. First public, timestamped articulation of AXI. |
| 1.1 (tentative) | 2026–2027 | First revision incorporating AXI Working Group contributions, contingent on community engagement. |

Future revisions will preserve earlier version numbers so products can cite conformance to a specific AXI version.

---

## 21. References


[1] TechCrunch. (2026, February 27). *ChatGPT reaches 900M weekly active users.* Based on OpenAI disclosures.

[2] Publicly reported annualized revenue and valuation figures for leading frontier AI companies (OpenAI, Anthropic), as of Q1 2026. Sources include TechCrunch, Fortune, Epoch AI data insights, and company disclosures.

[3] Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). *Trust, attitudes and use of artificial intelligence: A global study 2025.* KPMG & University of Melbourne (n = 48,340 across 47 countries).

[4] MIT Project NANDA (Challapally, A., lead author). (2025). *The GenAI Divide: State of AI in Business 2025.* MIT Media Lab.

[5] Gartner. (2025, June 25). *Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027.* Press release.

[6] Forrester Consulting. (2022). *The total economic impact of improved chatbot experiences.* And subsequent consumer-experience research on chatbot purchase-abandonment rates, 2022–2025.

[7] Pew Research Center. (2025, April 3). *How the US public and AI experts view artificial intelligence.* (Public n = 5,410, AI expert n = 1,013.)

[8] Stanford HAI. (2025). *AI Index Report 2025.* Human-Centered Artificial Intelligence, Stanford University.

[9] Grand View Research. (2025). *Artificial intelligence market size & outlook 2024–2030*; *Conversational AI market size & outlook.*

[10] MarketsandMarkets. (2025). *Agentic AI market — global forecast to 2032* (and updated June 2025 report with revised figures).

[11] Market Research Future. (2025). *Digital human (AI avatars) market research report, 2035.* (And disagreeing forecasts from The Business Research Company, Data Bridge Market Research, and Gartner.)

[12] Goldman Sachs Research. (2024). *Humanoid robots: TAM and trajectory analysis.*

[13] Gartner. (2022, August 31). *Forecast of contact-center labor cost displacement by conversational AI through 2026.*

[14] Digital Commerce 360 reporting on Gartner forecast (October 21, 2025): AI agents and $15 trillion in B2B purchases by 2028.

[15] Gartner. (2024, September 9). *Multimodal generative AI forecast: 40% of GenAI solutions by 2027.*

[16] New Market Pitch. (2025). *Humanoid robotics funding tracker, 2025.* Also Yole Group's cumulative humanoid funding analysis, October 2025.

[17] McKinsey & Company. (2025). *The state of AI: How organizations are rewiring to capture value* (n = 1,993 executives, 105 countries).

[18] S&P Global Market Intelligence. (2025). *AI project abandonment trends, 2024–2025.*

[19] Forrester Consulting. (2022). *The total economic impact of improved chatbot experiences.*

[20] Gartner. (2024, July 9). *64% of customers would prefer that companies didn't use AI for customer service.* Press release (n = 5,728).

[21] Norman, D. A. (2013). *The design of everyday things* (Revised and expanded ed.). Basic Books. (Original work published 1988.)

[22] Shneiderman, B. (1987). *Designing the user interface: Strategies for effective human-computer interaction* (1st ed.). Addison-Wesley.

[23] Nielsen, J. (1994). Enhancing the explanatory power of usability heuristics. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 152–158.

[24] Picard, R. W. (1997). *Affective computing.* MIT Press.

[25] Gunning, D. (2017). *Explainable artificial intelligence (XAI).* DARPA / I2O Technical Report.

[26] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115.

[27] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Paper 3, 1–13.

[28] Shneiderman, B. (2022). *Human-centered AI.* Oxford University Press.

[29] Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.

[30] Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4), 391–444.

[31] Bubeck, S., et al. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv:2303.12712.

[32] Yao, S., et al. (2023). ReAct: Synergizing reasoning and acting in language models. *ICLR 2023.* arXiv:2210.03629.

[33] Schick, T., et al. (2023). Toolformer: Language models can teach themselves to use tools. *NeurIPS 36.*

[34] Park, J. S., et al. (2023). Generative agents: Interactive simulacra of human behavior. *UIST 2023.*

[35] Schluntz, E., & Zhang, B. (2024, December 19). *Building effective agents.* Anthropic Engineering.

[36] Anthropic. (2024, October 22). *Introducing computer use, a new Claude 3.5 Sonnet, and Claude 3.5 Haiku.*

[37] OpenAI. (2025, January 23). *Introducing Operator.*

[38] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. *ACM SIGKDD 22*, 1135–1144.

[39] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *NeurIPS 30*, 4765–4774.

[40] Microsoft Aether. (2022). *Overreliance on AI: Literature review.*

[41] Radford, A., et al. (2021). Learning transferable visual models from natural language supervision. *ICML 38.*

[42] Alayrac, J.-B., et al. (2022). Flamingo: A visual language model for few-shot learning. *NeurIPS 35.*

[43] Liu, H., Li, C., Wu, Q., & Lee, Y. J. (2023). Visual instruction tuning. *NeurIPS 36.*

[44] OpenAI. (2023). *GPT-4V(ision) system card.*

[45] Gemini Team, Google. (2023). *Gemini: A family of highly capable multimodal models.* arXiv:2312.11805.

[46] Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1–3), 139–159.

[47] Liu, Y., et al. (2024). Aligning cyber space with physical world: A comprehensive survey on embodied AI. arXiv:2407.06886.

[48] Norman, D. A. (2004). *Emotional design: Why we love (or hate) everyday things.* Basic Books.

[49] Lombard, M., & Ditton, T. (1997). At the heart of it all: The concept of presence. *Journal of Computer-Mediated Communication*, 3(2).

[50] Slater, M., & Wilbur, S. (1997). A framework for immersive virtual environments (FIVE). *Presence: Teleoperators and Virtual Environments*, 6(6), 603–616.

[51] Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience.* Harper & Row.

[52] Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80.

[53] Nielsen Norman Group. (2024). *AI-era interaction patterns: Accordion editing and apple picking.*

[54] Weizenbaum, J. (1966). ELIZA — a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45.

[55] Breazeal, C. (2002). *Designing sociable robots.* MIT Press.

[56] Ayers, J. W., Poliak, A., Dredze, M., et al. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 183(6), 589–596. (Replicated in Ruben et al., *Journal of General Internal Medicine*, 2025.)

[57] Google PAIR. (2019, continuously updated). *People + AI guidebook.*

[58] Suleyman, M. (2025). Published commentary on “Seemingly Conscious AI” and related warnings about systems that produce false impressions of consciousness.

[59] Vectara. (2026). *Hallucination leaderboard (HHEM benchmark).*

[60] Bailenson, J. N. (2021). Nonverbal overload: A theoretical argument for the causes of Zoom fatigue. *Technology, Mind, and Behavior*, 2(1).

[61] Xu, W. (2024). A “User Experience 3.0 (UX 3.0)” paradigm framework: User experience design for human-centered AI systems. arXiv:2403.01609.

[62] Industry writing on “Agent Experience (AX)” as a distinct discipline for agent-facing AI design, 2025–2026. (Kore.ai, DataGroomr, and related practitioner publications.)

[63] Biānjiè Research Group. (2025). *What is Artificial Experience (AX)?* Industry essay positioning AI maturation as artificial experience.

[64] Industry writing on “Presence Continuity Layer” proposals for AI identity and memory infrastructure, 2026.

### Author’s Prior Work

[A1] Yeluri S. D. S. Sri Vardhan. (2024). Effects of business intelligence tools on financial performance of IT industry. *AIP Conference Proceedings*, 2971, 020036. https://doi.org/10.1063/5.0196170

(Note: The prior work above is in an adjacent field — business intelligence in IT — and is cited here for completeness. It does not establish domain expertise in HCI or AI, and the reader should weigh this paper on its own merits.)

---

## 22. Acknowledgments

The author thanks early readers of draft versions of this paper for their feedback, the open research communities of HCI, HCAI, and AI whose published work made this synthesis possible, and the Sripto Corporation team for supporting the time required to prepare this document. Any errors, omissions, or misjudgments are solely the author’s.

---

## 23. About the Author and Sripto Corporation

**Yeluri S. D. S. Sri Vardhan** is the Founder and Chief Executive Officer of Sripto Corporation Private Limited. He is the originator of the AXI framework proposal and the principal author of this whitepaper. His prior peer-reviewed publication, *Effects of Business Intelligence Tools on Financial Performance of IT Industry*, appeared in *AIP Conference Proceedings*, Volume 2971, 020036 (2024) [A1] — a paper in an adjacent field that is not directly related to the subject matter of this whitepaper.

His current work centers on the architecture of AI systems that are experienced by users as present, embodied, comfortable, and trustworthy — on the engineering of the bridge between artificial intelligence and the humans who use it.

Correspondence regarding the AXI framework, the AXI Working Group, and the AXI Technical Report Series may be directed to Sripto Corporation Private Limited.

**Sripto Corporation Private Limited** is an Andhra Pradesh, India-based company building AI products oriented around the experience layer. Sripto’s flagship product, *Let Me Teach* (2026), is an interactive explainer that teaches any topic in real time. Sripto’s products are designed against the AXI Stack, Principles, and Metrics described in this whitepaper.

**Conflict of interest.** This whitepaper is written by the CEO of the company that builds *Let Me Teach*. This is a material conflict of interest, disclosed at the top of this document and again in Section 11 and Section 13.

---

## 24. How to Cite This Paper

**Chicago / Author-Date:** Yeluri S. D. S. Sri Vardhan. 2026. *Artificial eXperience Intelligence (AXI): Toward a Framework for the Experience Layer of Artificial Intelligence.* Version 1.0. Sripto Corporation Private Limited.

**APA 7:** Yeluri S. D. S. Sri Vardhan. (2026). *Artificial eXperience Intelligence (AXI): Toward a framework for the experience layer of artificial intelligence* (Version 1.0). Sripto Corporation Private Limited.

**IEEE:** Yeluri S. D. S. Sri Vardhan, “Artificial eXperience Intelligence (AXI): Toward a Framework for the Experience Layer of Artificial Intelligence,” v1.0, Sripto Corporation Private Limited, Apr. 2026.

**BibTeX:**

```
@techreport{yeluri2026axi,
  author       = {Yeluri S. D. S. Sri Vardhan},
  title        = {Artificial eXperience Intelligence (AXI): Toward a Framework for the Experience Layer of Artificial
Intelligence},
  institution  = {Sripto Corporation Private Limited},
  year         = {2026},
  month        = {April},
  type         = {Foundational Whitepaper},
  version      = {1.0},
  license      = {CC BY 4.0},
  url          = {https://axi.sripto.tech}
}
```

---