

The Human-Centered AI Literacy Course

Building Judgment, Agency, and Trust in AI-Mediated Work

— *A course in epistemic self-trust under technological mediation*

A guided learning experience in human-centered AI literacy.

I. Course Overview

This course develops AI literacy by focusing on the one thing no system can replace: your judgment.

Rather than teaching tools, the course helps you understand how AI generates language and how fluent outputs can subtly shape human thinking.

It teaches the conditions under which AI's relational effects on meaning, responsibility, and trust become visible and interpretable, enabling clearer judgment between confidence and grounding, suggestion and reasoning, assistance and authority.

Across the course, you'll practice calibrating trust and setting boundaries while staying authorial when working with AI in individual and collaborative contexts.

The course is designed to be accessible and grounded in how people actually learn, with no technical background required.

By the end, you'll leave with a clarified lens for interpreting AI-mediated work that reflects your values, preserves your judgment, and remains durable across future tools, roles, and contexts.

II. Why This Course Exists

As AI systems become more fluent, persuasive, and embedded in everyday work, learning, and decision-making, the central challenge is no longer technical proficiency. It is **human judgment**.

This course was designed to address a growing gap:

how people remain grounded, discerning, and ethically clear while working alongside systems that can generate confident language without understanding and coherence without accountability.

Rather than focusing on tools or interfaces, the course develops the **human capacities** that determine whether AI strengthens thinking or subtly distorts it.

III. What This Course Does

The Human-Centered AI Literacy Course equips learners to:

- Understand how modern AI systems generate outputs using accurate, non-technical mental models
- Interpret fluent AI language without mistaking confidence for reliability, calibrating trust deliberately rather than reflexively
- Maintain authorship and responsibility in AI-mediated work
- Navigate ambiguity with steadiness and care
- Collaborate with others in ways that protect collective intelligence
- Develop an interpretive stance toward AI grounded in values, authorship, and judgment that is adaptable as contexts and technologies change

The emphasis is on **durable interpretive skill**, not platform-specific knowledge.

IV. Who This Course Is For

The course is designed for adult learners across education, work, and civic settings, including:

- Students and educators
- Knowledge workers and professionals
- Organizational leaders and teams
- Institutions seeking responsible AI capacity-building

The course is especially valuable for learners who want to work with AI thoughtfully without outsourcing judgment, identity, or responsibility to the system.

No technical background is required. Learners can enter at any level of AI familiarity.

V. What Makes This Course Distinctive

This course centers human cognition, agency, and relational judgment throughout.

Key features include:

- **Human-centered mental models**
Learners understand what AI systems are doing—and what they are not doing—without being misled about intelligence or intent.

- **Cognitive and emotional literacy**
The course addresses how AI shapes reasoning, confidence, framing, and emotional response, using observational rather than anthropomorphic language.
- **A steady, non-alarmist tone**
Learners are supported in developing clarity and confidence, not fear or dependency.
- **Durability by design**
All concepts remain applicable as tools, models, and interfaces evolve.
- **Agency-preserving pedagogy**
Learning activities are reflective and practice-oriented, emphasizing authorship and responsibility rather than performance.

VI. Course Structure & Learning Arc

The course is organized into **seven core learning modules** and one final synthesis module, intentionally sequenced to support cognitive clarity, emotional safety, and durable judgment:

1. **Foundations** — Functional Fluency & Early Metacognition
2. **How AI Thinks** — Cognitive Partnership and Mental Models
3. **Bias** — Meaning, Authority, and Distortion in AI-Mediated Work
4. **Boundary Work** — Self-Regulation, Over-Reliance, and Cognitive Integrity in AI-Supported Work
5. **Emotional Literacy** — Working With a Fluent, Non-Human Partner
6. **Navigating Ambiguity** — Calm Decision-Making Under AI-Generated Uncertainty
7. **Working With Others** — Collaboration and Collective Intelligence in AI-Mediated Environments
8. **Final Synthesis — Lens Audit** — A structured reflection on how values, identity, and judgment now shape one’s interpretation of AI-mediated work.

Rather than ending in a fixed framework, the course culminates in recognition: how your judgment and values now show up when AI is present—without requiring a single finalized or defensible way of working.

In addition to the core course, learners may explore an optional bonus module on Relational AI & Shared Sense-Making. This module is exploratory, non-credentialed, and designed as an interpretive extension rather than a continuation of the course.

VII. Pedagogical Integrity & Rigor

The course is rigorous without being performative.

- Cognitive load is carefully paced, with stabilization preceding destabilizing insights
- Assignments are ungraded or completion-based, designed to slow thinking and support reflection
- Metacognition is explicit and recurring
- Emotional safety is treated as instructional infrastructure, especially in identity-relevant modules

This approach makes the course suitable for diverse learners while preserving depth and seriousness.

VIII. Outcomes & Impact

Learners complete the course with:

- A clear, accurate mental model of AI behavior that supports interpretation without over- or under-trust
- Greater confidence in their own judgment
- Practical habits for calibration and boundary-setting
- Emotional steadiness when AI feels persuasive or uncertain
- Stronger norms for responsible collaboration
- A durable personal stance toward AI use marked by epistemic self-trust, interpretive clarity, and retained authorship.

These outcomes position learners as stewards of judgment, not just users of technology.

IX. Why This Matters

For internal leadership, this course establishes a values-aligned foundation for AI engagement that protects people, institutional credibility, and long-term capacity.

For funders and partners, it offers durable infrastructure for human agency in an AI-shaped

future—aligned with responsible innovation, workforce resilience, education, and ethical technology adoption.

X. In One Sentence

The Human-Centered AI Literacy Course equips people to think clearly, act responsibly, and remain fully human in environments where intelligence is increasingly shared with machines.