

# The Human-Centered AI Literacy Course

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

*A Course Brief for External Readers*

---

### 1. The Moment This Course Responds To

Artificial intelligence systems are now woven into everyday decision-making. They summarize information, generate options, frame problems, and shape what feels reasonable or complete. As these systems become more fluent and accessible, many people are gaining functional competence with AI tools faster than they are developing the judgment required to use them well.

This creates a subtle but consequential gap. Coherent outputs can feel authoritative even when they compress uncertainty, smooth over disagreement, or obscure context. Responsibility, however, does not shift with fluency. When judgment is displaced rather than supported, decision quality degrades quietly. Often, this happens without anyone noticing until consequences surface.

The Human-Centered AI Literacy Course exists to address this gap not by teaching people how to use AI more efficiently, but by strengthening the human capacities that matter when AI systems are already present.

---

### 2. What the Course Is (and Is Not)

The Human-Centered AI Literacy Course is a human-centered approach to AI literacy, focused on judgment, agency, and responsibility in AI-shaped decision environments.

**It is not a course about mastering tools.**

**It is a course about remaining accountable when tools are fluent.**

Accordingly, the course is **not**:

- tool training or software instruction
- prompt optimization or productivity coaching

- mastery-oriented or skills-first education

It is:

- an orientation to how AI systems shape decision environments
- a framework for recognizing the difference between fluency and understanding
- a structured space for developing discernment, boundary awareness, and responsible use

Rather than helping learners “get better at AI,” the Human-Centered AI Literacy Course helps people remain grounded, reflective, and accountable while working alongside increasingly persuasive systems.

---

### 3. The Core Distinction: Fluency Is Not Understanding

At the heart of the course is a simple but often overlooked distinction: **fluency is not understanding**.

AI systems can produce outputs that are coherent, confident, and persuasive without any comprehension of meaning, stakes, or consequence. Fluency is a property of generated output, while understanding is a human capacity tied to context, interpretation, and responsibility.

This distinction matters because fluent outputs can feel like understanding, particularly when summarizing documents, conversations, or complex decisions. In those moments, disagreement, uncertainty, and nuance are often removed, yet these are precisely the elements that judgment depends on. The course equips learners to recognize this dynamic as it happens, so coherence is not mistaken for care or authority.

---

### 4. How the Course Approaches Learning

This distinction shapes how the course is designed.

The Human-Centered AI Literacy Course is organized around **orientation rather than optimization**. Instead of progressing through tools, techniques, or checklists, learners develop a set of interrelated lenses for interpreting how AI systems influence attention, framing, confidence, and responsibility.

Learning is cumulative rather than linear. The emphasis is on recognition before response.

Where many AI learning experiences focus on what to do, this course focuses on how decision environments are being shaped and what remains a human obligation within them.

Its learning posture is guided by three consistent principles:

- attention before action
- interpretation before reliance
- accountability as a human responsibility

Throughout the course, learners practice noticing how AI-shaped decision environments are altered by summaries, suggestions, defaults, and apparent authority, before deciding how and whether to act.

---

## **5. Who the Course Is For**

Judgment is a human and civic capacity, not a specialized technical skill.

The Human-Centered AI Literacy Course is designed for anyone navigating AI-shaped decision environments, including:

- educators and students
- professionals across sectors
- leaders, collaborators, and practitioners
- individuals encountering AI systems in everyday work

No technical background is assumed, and the course does not gatekeep participation. It supports thoughtful engagement wherever AI systems influence interpretation and choice.

---

## **6. Why This Matters Beyond the Classroom**

As AI systems become embedded across institutions and workflows, the risks they introduce are rarely dramatic or malicious. More often, they are gradual: premature trust, reduced deliberation, and the normalization of fluent outputs as decision shortcuts.

By strengthening interpretive awareness and boundary-setting, the Human-Centered AI Literacy Course contributes to healthier decision cultures that preserve accountability, resist automation bias, and maintain space for human judgment even as systems become more capable.

The value of the course lies not in short-term efficiency gains, but in sustaining the conditions for responsible action over time.

---

## 7. How to Understand This Course

The Human-Centered AI Literacy Course is best understood as an orientation.

It does not offer prescriptions, guarantees, or mastery. Instead, it supports people in moments where **judgment cannot be fully outsourced**, even as AI systems participate in the decision environment.

A concise way to describe the course is this:

**The Human-Centered AI Literacy Course is a structured approach to strengthening human judgment, agency, and responsibility in AI-shaped contexts.**

---

*The Human-Centered AI Literacy Course*

Designed by Matthew Agustin