

# AI Ethics in Practice

## Companion Guide to the AI Ethics Rubric for L&D

Real-world examples and reflection prompts to  
bring ethical AI principles to life in L&D contexts



# How to Use This Guide

Ethical decision-making in AI isn't just about avoiding harm—it's about intentionally designing and shaping AI systems for trust, equity, and human dignity from the start. *AI Ethics in Practice* is your bridge between principles and action. As a companion to the *AI Ethics Rubric*, it offers concrete examples and guiding questions, making it easier to facilitate team discussions, guide procurement decisions, interact with vendors, audit AI usage, and foster governance.

---

## What's Inside

- 10 Ethical Dimensions in Practice – one section for each dimension in the AI Ethics Rubric
- Real-world examples – to illustrate common dilemmas and design trade-offs
- Reflection questions – to guide team discussions or self-assessment

## Ways to Use It

- In project planning – Integrate relevant sections into your design briefs or kickoff workshops
- During procurement – Use the questions to assess vendor practices and tool capabilities
- As part of reviews – Facilitate team discussions using examples to surface ethical considerations early
- For capability-building – Pair with the upcoming Job Aid & Worksheets to train staff in responsible AI use

## Our Recommendations

- Start with one dimension – You don't need to cover all ten at once; focus where the stakes or risks are highest
- Document your thinking – Record your answers, decisions, and trade-offs for transparency and future review
- Adapt to your context – The prompts are meant to be tailored; reword, combine, or expand to fit your culture and workflow



# Purpose & Value Alignment


Looks at whether the use of AI is intentional, pedagogically sound, and aligned with the organizational core values and strategic goals. It emphasizes using AI to enhance—not replace—human-centered learning experiences and ensures that deployments are driven by meaningful learning outcomes rather than novelty or convenience.

## Example

An organization pilots AI agents to support learners with real-time guidance during onboarding and training. As the agents become more autonomous—able to suggest content, monitor learner behavior, and escalate issues—the L&D team pauses to ask: Do these capabilities reflect our values of learner autonomy, trust, and growth? Upon review, they disable default features that flag "underperformance" or track user activity without consent. Instead, they re-design the agents to act as optional learning companions—offering help when asked, surfacing resources based on learner goals, and explaining how suggestions are generated.

## Questions to Ask

- Does this AI tool support—not contradict—our mission, learning philosophy, or organizational culture?
- What assumptions about productivity, efficiency, or control are baked into this tool?
- Have we considered how the tool might shape (or shift) learner or educator behaviors?



# Transparency & Explainability

Assesses whether learners and stakeholders are communicated with and understand what the AI does, how it works, and why it makes specific decisions or recommendations.

## Example

A course includes an AI-powered recommendation engine that suggests personalized learning paths. When a learner clicks on a suggestion, a pop-up explains:

*“This recommendation was generated by an AI tool using your past course completions and assessment scores. No personal identity data was used.”*

Additionally, a help link offers more detail on how the algorithm makes decisions.

## Questions to Ask

- Do learners or educators understand what the AI is doing and why?
- What happens when an AI tool makes a decision learners disagree with—can they trace the logic?
- Have we given learners just enough explanation to build trust without overwhelming?





# Bias & Fairness


Focuses on identifying and addressing bias in AI systems—whether embedded in training data, algorithms, or design choices—to prevent the marginalization or exclusion of any learner group. Fairness ensures that AI tools provide equitable access, representation, and outcomes for all learners, regardless of background, ability, or identity.

## Example

An adaptive learning platform uses AI to recommend professional development courses. Over time, the system disproportionately recommends technical tracks to male employees and soft skills to female employees. An audit reveals that historical enrollment data reinforced existing workplace stereotypes. The algorithm is adjusted to remove gender as a proxy variable and a fairness review is introduced to monitor future recommendation patterns.

## Questions to Ask

- Do we check whether AI outputs reinforce stereotypes or exclude certain groups?
- Whose data and norms are shaping the model's behavior?
- What's our process for regularly auditing outputs for bias—qualitative and quantitative?



# Data Ethics & Privacy

Ensures responsible handling of learner data—collecting only what’s needed, securing it properly, and being transparent about how it’s used , stored, and archived.

## Example

In a talent development platform, data is collected to support learning analytics—but only the information necessary to provide meaningful feedback and improve the learning experience is captured. No extraneous personal or behavioral data is gathered. All learner data is anonymized, used solely for instructional improvement, and never factored into performance evaluations or HR decisions. The system adheres to a strict data retention and deletion policy, with transparent guidelines for data archival, access, and removal.

## Questions to Ask

- Are we collecting only the data we need—and can we justify why we need it?
- Is learner data anonymized, protected, and governed by a clear deletion policy?
- How are we communicating our data practices to learners and other stakeholders?



# Accountability & Oversight

Defines who owns ethical oversight for AI tools. This includes governance mechanisms, review cycles, and documentation of risks and remediation plans.

## Example

A university's instructional design department establishes an AI Use Review Committee to evaluate all proposed AI tools for learning content development. The committee includes faculty, students, IT specialists, accessibility advisors, and ethics representatives to ensure a range of perspectives. It meets monthly to assess potential ethical risks, document decisions, and track approvals or rejections. All outcomes are logged in a centralized dashboard and published on the intranet, ensuring traceability and institutional oversight.

## Questions to Ask

- Who is responsible for reviewing, approving, and updating our AI tools?
- Is there a clear, documented process for ethical review?
- How do we track and act on ethical concerns that arise after deployment?





# Learner Agency & Consent

Informs learners when AI is in use, obtains explicit consent where appropriate, and provides meaningful ways for them to influence how AI shapes their learning experience. This includes the ability to opt out, adjust preferences, and question or override AI-driven decisions that impact their development or progression.


## Example

Learners entering an AI-enabled coaching simulation are informed that their input will shape the conversation and that the AI will adapt responses based on their choices. They are offered a clear opt-out option, with the choice to engage in a human-facilitated coaching path instead—without any impact on their performance review or learning record. Learners who opt out can rejoin the AI simulation at any time, and those who continue are able to reset or reroute the conversation whenever they choose.

## Questions to Ask

- How are learners informed that AI is in use? Is the language understandable and non-technical?
- What options do learners have to opt out, adjust preferences, or challenge AI-generated outcomes?
- Are there any implicit pressures to engage with the AI tool?





# Ethical Use of GenAI

Encompasses issues of copyright, authorship, accuracy, creativity, and responsible use of large language models (LLMs). It includes practices such as fact-checking AI-generated content, maintaining citation standards, clarifying when and how content is AI-assisted, and ensuring appropriate human-AI collaboration.

## Example

An instructional designer uses ChatGPT to draft knowledge check questions for a cybersecurity course. During review, they discover that several AI-generated questions reference outdated practices and cite fictional tools not found in internal IT policy—a common example of GenAI “hallucination.” To ensure accuracy, each question is verified against current policy documents and other primary sources of information, corrected for relevance, and labeled as AI-assisted within the instructional file. The designer also adds human-written commentary to clarify intent and context, ensuring the final materials are factually sound.

## Questions to Ask

- Are we clearly labeling and curating AI-generated content before using it in learning?
- What’s our process for fact-checking, revising, or contextualizing generative AI outputs?
- Are we treating GenAI as a co-creator or as a shortcut?



# Vendor Due Diligence

Focuses on ethical procurement practice—including how vendors are evaluated through contracts, documentation, and alignment with your values. This dimension ensures that third-party tools and providers uphold responsible practices across design, data use, and sustainability.

## Example

When evaluating a text-to-voice AI tool, the learning team embeds ethical criteria into the RFP process, requiring vendors to disclose:

- Sources of training data
- Data privacy and retention practices
- Model update frequency and governance
- Environmental impact metrics

One vendor is disqualified for failing to provide sufficient documentation of GDPR compliance, underscoring the team's commitment to responsible procurement and regulatory alignment.

## Questions to Ask

- What ethical commitments or practices has the vendor made public (e.g. around data, accessibility, sustainability)?
- Can we review their documentation on model updates, privacy, or environmental impact?
- What is their roadmap for responsible AI use—not just current capabilities?



# Inclusive AI Design

Focuses on how AI tools must consider diverse learner needs—across languages, abilities, contexts, and cultures. Testing for accessibility and relevance is essential.

## Example

An L&D team uses an AI tool to generate text-based scenarios for a global leadership training program. During review, they notice the AI defaults to Western names, office settings, and idiomatic language that may confuse ESL learners. To improve accessibility and inclusion, the team adjusts prompts and retrains the model to reflect more global contexts, clearer language structures, and a wider range of names and roles. The revised scenarios are free of jargon, and better suited to learners with varied language backgrounds.

## Questions to Ask

- Does this tool reflect the diversity of our learner population—culturally, linguistically, and contextually?
- How are we involving underrepresented voices in the testing or design process?
- Are we designing for different language backgrounds, sub-cultures, prior knowledge/experience levels, and accessibility needs?





# Environmental Sustainability

Acknowledges the environmental impact of using AI tools in learning, including the energy consumption, carbon footprint, and resource intensity of training, hosting, and running AI systems.

## Example

When faced with a decision between two AI video creation platforms, the L&D team evaluates the vendors' environmental practices. Subsequently, they select the vendor whose technology runs on carbon-neutral cloud infrastructure and who publishes a sustainability report. To further reduce environmental impact, the team also develops internal guidelines to limit unnecessary re-renders of AI-generated video to reduce compute cycles.

## Questions to Ask

- What energy demands does this tool introduce (e.g. compute-heavy features)?
- Have we considered cloud provider sustainability, or ways to reduce unnecessary AI usage?
- Are we valuing environmental impact as well as performance and cost?