

AI Ethics Rubric for L&D – 10 Dimensions of Consideration

10 Dimensions to Guide Responsible AI Use in Learning Design

This rubric is designed for L&D professionals, instructional designers, learning strategists, educators, and anyone integrating AI into learning ecosystems. It offers a practical and holistic lens for assessing the ethical dimensions of AI tools and initiatives in context.

These 10 dimensions are not a checklist, but a design framework—intended to prompt reflection, support better questions, and guide more intentional choices as we navigate the complex intersection of learning and artificial intelligence.

Each dimension is accompanied by a clear definition and can be explored independently or collectively, depending on where you are in your AI journey.



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.



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.



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.



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.



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.



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.



Accountability & Oversight

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



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.



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.



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.

* Generative AI (GenAI) refers to a specific type of AI that can create new content—such as text, images, code, or audio—based on patterns in its training data. Unlike traditional AI systems that classify or predict, GenAI tools such as ChatGPT, Claude, DeepSeek, or Gemini generate original outputs, which raises unique ethical considerations for learning design.

Using the AI Ethics Rubric as a Design Tool

The **AI Ethics Rubric for L&D** isn't just a checklist—it's a *design tool* to help organizations and teams navigate complexity, reflect critically, and take responsible action as they implement AI in learning environments.

Traditionally, ethics has been framed as a matter of judgment: Is this right or wrong? But in practice—especially in L&D—ethical decisions unfold more like design challenges than binary dilemmas. They're messy, contextual, evolving, and rarely solved with a single right answer.

Inspired by the concept of [Ethics as Design](#) by Caroline Whitbeck, this rubric invites you to treat ethics not as a static set of rules, but as a creative, iterative process:

- **Design solutions, don't just identify gaps** — Use each dimension to uncover not just what's missing, but how you might redesign or improve the approach. Brainstorm alternatives, weigh tradeoffs, and co-create solutions with learners and stakeholders.
- **Embrace complexity** — Ethical issues rarely have one right answer. Use the dimensions to explore tensions (e.g., automation vs. agency, personalization vs. privacy) and find balanced, context-sensitive paths forward.
- **Design iteratively and revisit over time** — Like AI systems themselves, your ethical stance should evolve. Revisit this tool often. Adjust as new technologies, insights, and impacts emerge. Use the rubric at different stages: product procurement, design, implementation, and review.

How to Begin:





- Pick the project or AI tool you're evaluating.
- Go through each of the 10 dimensions (if applicable, feel free to adapt the dimensions to your project as you see fit).
- Ask: *Where does this project land—Emerging, Developing, or Leading Practice?*
- Use your insights to refine, reframe, or re-align the initiative in ways that balance equity, trust, and purpose.







The AI Ethics Rubric

The AI Ethics Rubric for L&D is a practical reflection tool designed to help learning teams evaluate AI projects across 10 key dimensions. Each dimension is scored at one of three levels:

- **Emerging:** Minimal ethical consideration or unstructured practice
- **Developing:** Some safeguards or practices are in place, but inconsistently applied
- **Leading:** Ethical principles are fully integrated, proactive, and continuously improved

These levels are not about scoring organizations, but about supporting thoughtful assessment of projects or initiatives. Some teams may intentionally begin in the “Emerging” phase due to experimentation or capacity; others may aim for “Leading” as a long-term aspiration. The value lies in the intentional reflection and design choices, not the label itself. Use the rubric to reflect, ask better questions, design responsibly, and guide AI implementation that aligns with your values and your learners’ needs.

	Dimension	Emerging	Developing	Leading
1.	 Purpose & Value Alignment	AI is adopted without clear alignment to learning goals, organization’s values, or strategy.	AI tools are somewhat aligned with organizational as well as learning goals and values, but not consistently evaluated for fit or impact.	AI initiatives are purpose-driven, learner-centered, and fully aligned with organizational mission and goals.
2.	 Transparency & Explainability	AI decisions and processes are opaque to learners and stakeholders.	Basic explanations are provided, but transparency is inconsistent or overly technical.	AI operations are transparent, with clear, accessible explanations for outputs and recommendations.
3.	 Bias & Fairness	No formal process exists to identify or address bias in AI systems or data.	Initial efforts to reduce bias exist, such as limited testing or content review.	Bias is continuously monitored, with structured audits and inclusive design baked into AI development.
4.	 Data Ethics & Privacy	Data collection is excessive, ad hoc, or conducted at scale with little attention to consent, usage, or privacy protection.	Privacy policies are in place but not fully communicated or applied across use cases.	Data use is limited to as needed, consent-based, and learner-centered, with robust protection policies.

5.	 Accountability & Oversight	No designated oversight or governance for ethical AI implementation.	Some ethical review occurs, often informally or post-implementation.	Ethical governance is established with clear roles, accountability, and proactive monitoring.
6.	 Learner Agency & Consent	Learners are not informed about AI use and have no control or options.	Learners are informed about AI but have limited control or recourse.	Learners can opt in or out, modify how AI interacts with them, and challenge or override AI decisions.
7.	 Ethical Use of GenAI	GenAI is used without guidance on attribution, accuracy, or appropriateness.	Basic guidelines are in place for using GenAI, but practices vary widely.	GenAI use is transparent, fact-checked, attributed when appropriate, and reinforces creativity and originality.
8.	 Vendor Due Diligence	Vendors are selected based on functionality, brand recognition, or price, with no ethics review.	Vendors are asked about ethics, but due diligence is inconsistent, limited, or lacks depth.	Vendors are evaluated using pre-vetted ethical criteria from diverse stakeholder groups, including transparency, bias mitigation, and sustainability.
9.	 Inclusive AI Design	AI tools are not evaluated for accessibility, inclusivity, cultural context, or learner diversity.	Some inclusive design practices are adopted, but not across all tools or user groups.	AI tools are designed and tested for accessibility, inclusion, and equity from the start, and it is an ongoing evaluation practice.
10.	 Environmental Sustainability	No consideration is given to AI's energy use, opportunity cost, or environmental impact.	Environmental impact is acknowledged but not systematically assessed or tracked.	AI tools are chosen and configured with environmental impact in mind, favoring low-energy models and green hosting.